Identifying the Causal Effects of Cross-World Policies - Supplemental

Anonymous Author(s)
Affiliation
Address
email

Appendix

Proofs

Proposition 1. *For a CWPG $\mathcal{G}'(\mathcal{G}, \boldsymbol{Y}, \boldsymbol{A})$ built from $\mathcal{G}$ with observed variables $\boldsymbol{V}$, we have:*

1. *For any $V'_i \in \boldsymbol{V'}$, $|CP_{\mathcal{G}}(V'_i)| = 1$*
2. *For any $V'_i \in \boldsymbol{V'}$, $|Pa_{\mathcal{G}'}(V'_i)| = |Pa_{\mathcal{G}}(CP_{\mathcal{G}}(V'_i))|$*
3. *For any $V_i \in \boldsymbol{V}$ and any $V'_i \in CP_{\mathcal{G}'}(V_i)$ we have that $\bigcup_{V'_j \in Pa_{\mathcal{G}'}(V'_i)} CP_{\mathcal{G}}(V'_j) = Pa_{\mathcal{G}}(V_i)$*

Proof. (1) follows directly from the construction of a CWPG. (2) follows from steps 2c and 2d in the CWPG construction, which enforces that for any $V'_i \in \boldsymbol{V'}$ that has a copied version $V^c_i \in \boldsymbol{C}$, $\text{Ch}_{\mathcal{G}'}(V'_i) \cap \text{Ch}_{\mathcal{G}'}(V^c_i) = \emptyset$; so the addition of the copy nodes do not add any "extra" incoming edges to any variable in the graph. (3) follows from the construction of the CWPG, together with (1) and (2). □

The above Propositions give us the guarantee of a one to one function from the counterparts of the parents of some variable V'_i in $\mathcal{G}'$, and the parents of $\text{CP}_{\mathcal{G}}(V'_i)$, a fact that is important to keep in mind when proving functional equivalences.

Lemma 1. *Let $p(\boldsymbol{V})$ be a distribution nested Markov relative to an ADMG $\mathcal{G}$. Fix a set of outcome variables, $\boldsymbol{Y} \subset \boldsymbol{V}$, and treatment variables, $\boldsymbol{A} \subset \boldsymbol{V} \setminus \boldsymbol{Y}$, and build a CWPG $\mathcal{G}'(\mathcal{G}, \boldsymbol{Y}, \boldsymbol{A})$. Choose an assignment $\boldsymbol{v}$ for the variables in V, and let copy$(\boldsymbol{v})$ be an assignment that copies the assignments of $\boldsymbol{v}$ to the corresponding variables and copy variables in $\boldsymbol{V'} = \boldsymbol{V} \setminus \boldsymbol{C} \cup \boldsymbol{C}$. Then for all values of $\boldsymbol{v}$ we have:*

$$p(\boldsymbol{V} = \boldsymbol{v}) = p'(\boldsymbol{V'} = copy(\boldsymbol{v})) = \sum_{\boldsymbol{C}} p'(\boldsymbol{V'} \setminus \boldsymbol{C} = \boldsymbol{v}, \boldsymbol{C})$$

Proof. Since we assume a functional causal model, the settings of the variables are deterministic functions of the exogenous error variables $\boldsymbol{\epsilon}$. By construction of $\mathcal{G}'$ shares the same distribution over the error variables, $p(\boldsymbol{\epsilon})$. Set $\boldsymbol{\epsilon}$ to the same value in both $\mathcal{G}$ and $\mathcal{G}'$ models. Now pick an observed variable $V'_i \in \boldsymbol{V'} \setminus \boldsymbol{C}$ which has no observed parents; the existence of such a variable follows from the acyclicity of $\mathcal{G}'$. By construction, their also exists a variable V_i in $\mathcal{G}$ which has the same assignment function, $\mathcal{F}_i$ as V'_i. Since V_i and V'_i share the same assignment function that only depends on the fixed error terms, and the error terms are fixed to the same values in both graphs, it follows that $V_i = V'_i$ whenever the error variables have the same values for both models. Set both variables to the value they take, and add them to the set F of fixed variables. If V'_i has a copy in $\boldsymbol{C}$ then it also takes on the same value due to functional equivalence. Set it to this value and add it to F.

Submitted to 34th Conference on Neural Information Processing Systems (NeurIPS 2020). Do not distribute.

Now pick a new observed variable $V'_j \in \boldsymbol{V'} \setminus \boldsymbol{C}$ which either has no observed parents, or all of its parents are in F.

We can prove the existence of such a variable as follows: Assume no such variable exists, and not all variables in $\mathcal{G}'$ have been added to F. Then there must exist some variable V'_k in $\mathcal{G}'$ which has a parent not in F. This parent must also have a parent not in F. Continuing this recursion, we will eventually reach a node with either no parents, or with all parent in F (by the finiteness of our variable set and acyclicity), which contradicts the initial assumption. Thus such a variable must exist, or all variables in the graph are in F.

Thus the variable V'_j, as well as its analogue in $\mathcal{G}$, V_j, both have parents in F set to the same value. From the shared assignment function between V'_j and V_j, and the equal setting of error variables across both models, it follows that $V'_j = V_j$. We can add these variables to F, along with any copied version of V'_j, and repeat this process until all variables are in F.

From this it follows that whenever the error variables are set the same in $\mathcal{G}$ and $\mathcal{G}'$, then $\boldsymbol{V} = \boldsymbol{v}$ implies $\boldsymbol{V'} = \text{copy}(\boldsymbol{v})$. Since the distribution over error variables is the same in both models, it follows that $p(\boldsymbol{V} = \boldsymbol{v}) = p'(\boldsymbol{V'} = \text{copy}(\boldsymbol{v}))$. That $p'(\boldsymbol{V'} = \text{copy}(\boldsymbol{v})) = \sum_{\boldsymbol{C}} p'(\boldsymbol{V'} \setminus \boldsymbol{C} = \boldsymbol{v}, \boldsymbol{C})$ follows from the fact that the nodes in C must take on the same value as the node they are a copy of in $\boldsymbol{V}' \setminus \boldsymbol{C}$, so all other assignment have zero probability mass in the observed distribution. □

Lemma 2. *For a graph $\mathcal{G}$ and its CWPG $\mathcal{G}'$ with variables $\boldsymbol{V}'$, if an ETT path specific policy intervention, $\boldsymbol{f}_{\pi_{\boldsymbol{A},\boldsymbol{Y}}}$, on $\mathcal{G}$ for treatments $\boldsymbol{A}$ and outcomes $\boldsymbol{Y}$ is expressible as an edge specific policy intervention $\boldsymbol{f}_\alpha$, we have:*

$$p(\boldsymbol{Y}(\boldsymbol{f}_\alpha)) = p'(\boldsymbol{Y'}(\boldsymbol{f}'_{\alpha'}))$$

where $\boldsymbol{Y}'$ are the counterparts of $\boldsymbol{Y}$ in $\mathcal{G}'$, α is the set of all edges in $\mathcal{G}$ starting a path in $\pi_{\boldsymbol{X},\boldsymbol{Y}}$, α' is the set $\{(A'X)'_{\rightarrow} | X \in \boldsymbol{C} \cup \boldsymbol{Y'}\}$ in $\mathcal{G}'$, and $f_A^{(AX)\rightarrow} \in \boldsymbol{f}$ is functionally equivalent to $f_{A'}^{(A'C)'_{\rightarrow}} \in \boldsymbol{f'}$ when $CP_{\mathcal{G}'}(A) = A'$

Proof. Under the recursive substitution definition for counterfactuals, $\boldsymbol{Y'}(f_{\alpha'})$ is equal to:

$$Y'(f_{\alpha'}) = Y'(\{f_{A'_i}^{(A'_iY')\rightarrow}(W_i) | A'_i \in \text{Pa}_{\mathcal{G}'}(Y') \cap \boldsymbol{A'}\}, \{V_i(f_{\alpha'}) | V_i \in \text{pa}_{\mathcal{G}}(Y') \setminus \boldsymbol{A'}\})$$

Similarly, the counterfactual $Y(f_{\alpha^*})$ is defined as:

$$Y(f_\alpha) = Y(\{f_{A_i}^{(A_iY)\rightarrow}(W_i) | A_i \in \text{Pa}_{\mathcal{G}}(Y) \cap \boldsymbol{A}\}, \{V_i(f_\alpha) | V_i \in \text{pa}_{\mathcal{G}}(Y) \setminus \boldsymbol{A}\})$$

From Proposition 1, and step (d) in the CWPG construction, one can see that $|\text{Pa}_{\mathcal{G}}(Y_i)| = |\text{Pa}_{\mathcal{G}'}(Y'_i)|$, and that each parent of Y_i has a one counterpart in $\text{Pa}_{\mathcal{G}'}(Y'_i)$. As such, to complete the proof we will show that the terms appearing in the recursive substitution definition of $Y'(f_{\alpha'})$ are functionally equivalent to their counterparts in $Y(f_\alpha)$. In $\mathcal{G}'$, the parents of Y'_i may be of one of three kinds of variables: a treatment variable in $\boldsymbol{A}'$, a copy variable in $\boldsymbol{C}$, or an outcome variable in $\boldsymbol{Y}'$. Since the elements of $\boldsymbol{Y'}$ are fixed for the purpose of evaluating a probability, the parents of Y'_i that are elements of $\boldsymbol{Y'}$ are vacuously equivalent to their counterparts in $\mathcal{G}$.

For the variables in A', we have by construction that they are set according to policy functions $f_{A'_i}^{(A'_iY)\rightarrow}(W'_i)$, and their counterparts which are parents of Y in $\mathcal{G}$ are set by equivalent functions $f_{A_i}^{(A_iY)\rightarrow}(W_i)$. Thus, the assignments to parents belonging in the treatment set will be identical across graphs as long as their inputs, W_i and W'_i, are the same.

We now look at variables in $\boldsymbol{C}$ which are parents of $\boldsymbol{Y}'$. By recursive substitution, all variables $C \in \text{Pa}(Y') \cap \boldsymbol{C}$ are set to $C(f_{\alpha'})$. For some C, if it does not have parents, then it is functionally equivalent to its counterpart in $\mathcal{G}$ (since its only dependence is on the exogenous noise terms which are the same across graphs). If C has parents its parents are either: (1) a variable $A'_i \in \boldsymbol{A'}$ whose value is set by a policy function $f_{A'_i}^{(A'_iY)\rightarrow}(W'_i)$, (2) an element of Y' which is fixed and identical across graphs, or (3) another variable in $\boldsymbol{C}$. This nesting ends only in the case of no parents, or in the case of (1) or (2). Due to acyclicity and finiteness, all terms in the recursive substitution definition of $C(f_{\alpha'})$ will eventually reach either the no parents case, case (1), or (2). Since assignments in the no parents case and in case (2) are functionally equivalent across graphs, we will have that $C(f_{\alpha'})$ is functionally equivalent to its counterpart in $\mathcal{G}$ as long as policy assignments to variables in $\boldsymbol{A'}$ are equivalent to there counterparts. Since the counterpart for a treatment variable shares an assignment function $f_{A'_i} = f_{A_i}$, this will be the case if we can show that the inputs to these functions are the same across graphs. If we show this, then the parents of Y' are functionally equivalent to their counterparts in $\mathcal{G}$ under the edge policy intervention, and the counterfactuals $Y_i(f_\alpha)$ and $Y'_i(f_{\alpha'})$ are thus functionally equivalent.

Recall that the inputs to a policy may either be the original, naturally occurring value of the variable V_i, or the counterfactual version of it under the policy $V_i(f)$. In $\mathcal{G}'$ under our assumed edge intervention, the former is represented explicitly by variables in $\boldsymbol{V'} \setminus \boldsymbol{C}$, while the later is represented explicitly by variables in $\boldsymbol{C}$. If the input corresponds to the naturally occurring value of the variable, then the input is functionally equivalent across graphs by the proof of Lemma 1.

If the input is a counterfactual, then look at a treatment variable A'_i in $\mathcal{G}'$ which has no other treatment variable as an ancestor (that such a variable exists comes from the acyclicity of the graph). The inputs to the policy function that assigns treatments to A'_i must all be natural valued variables, or counterfactuals that are equivalent to the natural value of the variable. This follows from rule 3 of the PO calculus (Malinsky et al., 2020). Thus, the output of the policy function on A'_i for edges going to nodes in $\boldsymbol{C}$ is functionally equal the output of the policy function on its counterpart in $\mathcal{G}$. We can now reuse the proof of Lemma 1 to pick a new variable A'_j which, if it has a treatment variable as an ancestor, it is a treatment variable whose policy inputs have been shown equivalent across graphs, which proves input equivalence across graphs for the assignment policy on A'_j. From this, input equivalence for all treatment assignment policies follows.

Thus, $Y_i(f_\alpha)$ and $Y'_i(f_{\alpha'})$ are functionally equivalent for any Y_i and counterpart Y'_i. From the equivalence of the error terms and their respective distribution, we thus have that

$$p(\boldsymbol{Y}(f_\alpha)) = p'(\boldsymbol{Y'}(f_{\alpha'}))$$

□

Theorem 1. *Let α' be the set of edges $\{(AX)_\rightarrow | A \in \boldsymbol{A'}; X \in \boldsymbol{C} \cup \boldsymbol{Y'}\}$ in $\mathcal{G}'$ and let $\boldsymbol{f}'_{\alpha'}$ be an edge policy intervention in $\mathcal{G}'$ on edges in α'. Assume that $\boldsymbol{f}_{\pi_{\boldsymbol{A},\boldsymbol{Y}}}$ is expressible as an edge intervention $\boldsymbol{f}_\alpha$ and $f_{A_i} \in \boldsymbol{f}$ is functionally equal to $f'_{A'_i} \in \boldsymbol{f'}$ when $CP_{\mathcal{G}'}(A_i) = A'_i$. Then $\boldsymbol{f}_\alpha$ is identified if and only if $\boldsymbol{f}'_{\alpha'}$ is.*

Proof. Assume that $\boldsymbol{f}_\alpha$ is not identified. Then there exists two parameterizations M_1 and M_2 for $\mathcal{G}$ such that $p_1(\boldsymbol{V}) = p_2(\boldsymbol{V})$ but $p_1(\boldsymbol{Y}(\boldsymbol{f}_\alpha)) \neq p_2(\boldsymbol{Y}(\boldsymbol{f}_\alpha))$, where p_i is the distribution under model M_i. Create a CWPG $\mathcal{G}'$ with variables $\boldsymbol{V'}$. Then by Lemma 1, we have $p'_1(\boldsymbol{V'}) = p'_2(\boldsymbol{V'})$. By Lemma 2 we have $p'_1(\boldsymbol{Y'}(\boldsymbol{f}'_{\alpha'})) = p_1(\boldsymbol{Y}(\boldsymbol{f}_\alpha))$ and $p'_2(\boldsymbol{Y'}(\boldsymbol{f}'_{\alpha'})) = p_2(\boldsymbol{Y}(\boldsymbol{f}_\alpha))$ and thus $p'_1(\boldsymbol{Y'}(\boldsymbol{f}'_{\alpha'})) \neq p'_2(\boldsymbol{Y'}(\boldsymbol{f}'_{\alpha'}))$. Thus $\boldsymbol{f}'_{\alpha'}$ is not identified.

Now assume that $\boldsymbol{f}'_{\alpha'}$ is not identified. Then follow a similar proof above to show that this implies $\boldsymbol{f}_\alpha$ is not identified. From this, our result follows. □

Theorem 2. *The Cross World Policy response $p(Y(\boldsymbol{f}_{\pi_{\boldsymbol{A},\boldsymbol{Y}}}))$ is identified from $p(\boldsymbol{V})$ if and only if the following hold:*

- *$\boldsymbol{f}_{\pi_{\boldsymbol{A},\boldsymbol{Y}}}$ is expressible as an edge intervention $\boldsymbol{f}_\alpha$ on $\mathcal{G}$*
- *$Ch_{\mathcal{G}'_{\boldsymbol{Y}^*}}(A'_i) \cap Dis_{\mathcal{G}'_{\boldsymbol{Y}^*}}(A'_i) = \emptyset$ for all $A'_i \in \boldsymbol{A'}$,*
- *No districts $D \in \mathcal{D}((\mathcal{G}'_{\boldsymbol{f}'_{\alpha'}})_{\boldsymbol{Y}^*})$ contain both a variable $V'_i \in \boldsymbol{V'} \setminus (\boldsymbol{C} \cup \boldsymbol{Y'})$ and a variable $V'_j \in \boldsymbol{C} \cup \boldsymbol{Y'}$*

If these hold, then the identifying formula is:

$$\sum_{\boldsymbol{Y}^* \setminus \boldsymbol{Y'}} \prod_{D \in \mathcal{D}(\mathcal{G}'_{\boldsymbol{Y}^*})} \phi_{\boldsymbol{V'} \setminus D}(p(\boldsymbol{V'}), \mathcal{G}')\Big|_{\{A'_i = f'_{A_i}(W'_i) | A'_i \in \boldsymbol{A'} \cap Pa^Y(D)\}} \tag{1}$$

Where $Pa^Y(D)$ are parents of D along the edges $\{(A'X)'_\rightarrow | A' \in \boldsymbol{A'}; X \in \boldsymbol{C} \cup \boldsymbol{Y'}\}$

Proof. By Theorem 1, proving identification for the Cross World Policy response $p(Y(\boldsymbol{f}_{\pi_{\boldsymbol{A},\boldsymbol{Y}}}))$ in $\mathcal{G}$ can be done by proving identification for $p'(Y'(\boldsymbol{f}'_{\alpha'}))$ on the graph $\mathcal{G}'$.

The first thing to note is the construction of the graph $\mathcal{G}'_{\boldsymbol{f}'_{\alpha'}}$ which represents the graph under the policy intervention and will determine the variables belonging in $\boldsymbol{Y}^*$. Contrary to usual constructions, no edges are removed in the creation of this graph, and some may be added. To see that this will always be the correct construction for $\mathcal{G}'_{\boldsymbol{f}'_{\alpha'}}$ note that (1) We can construe $\boldsymbol{f}'_{\alpha'}$ as additionally defining policies on edges $\{(A'X)'_{\rightarrow} | A' \in \boldsymbol{A}'; X \notin \boldsymbol{C} \cup \boldsymbol{Y}'\}$ which set A' to its natural value, and (2) we assume for all $f'_i \in \boldsymbol{f}'$ we have $W_i^{\mathfrak{M}} \neq \emptyset$. From this we can see that all nodes will have some edge with a policy which assigns the natural value. Since the inputs to this "policy" are simply the original inputs to A' in $\mathcal{G}'$, we do not remove any edges to A' in the original graph (though edges may be added).

Given this, we first show that the three conditions of Theorem 2 imply identification of $\boldsymbol{f}'_{\alpha'}$. Since our expression is originally a path intervention a necessary condition for identification is for the first point be true; that $\boldsymbol{f}_{\pi_{\boldsymbol{A},\boldsymbol{Y}}}$ be expressible as an edge intervention $\boldsymbol{f}_{\alpha}$ (Theorem 5.2 of Shpitser & Tchetgen (2016)). To prove that $p'(Y'(\boldsymbol{f}'_{\alpha'}))$ (and hence $p(Y(\boldsymbol{f}_{\alpha}))$) is identified, we use Theorem 3 from Shpitser & Sherman (2018), which implies identification for edge policies if (1) $\boldsymbol{Y}^*(\boldsymbol{A}' = \boldsymbol{a}')$ is identified for any assignment $\boldsymbol{a}'$, and (2) the edge assignments to the districts in $\mathcal{G}'_{\boldsymbol{f}'_{\alpha'}}$ are consistent (there exist no recanting districts).

If the second point of Theorem 2 is true, then this will imply that $\boldsymbol{Y}^*(\boldsymbol{A}')$ is identified, using the same observation used in Theorem 3 of Sani et al. (2020) (that Theorem 60 of Richardson et al. (2017), the one line ID algorithm, is valid even when $\boldsymbol{A}$ and $\boldsymbol{Y}$ are not disjoint).

If the third condition of Theorem 2 holds then it will follow that there will be no districts $D \in \mathcal{D}((\mathcal{G}'_{\boldsymbol{f}'_{\alpha'}})_{\boldsymbol{Y}^*})$ with inconsistent edge assignments. In our setting, only edges from some A' to an element of $\boldsymbol{C} \cup \boldsymbol{Y}'$ will have an assignment different from the natural value, hence, only the existence of a variable $V'_i \in \boldsymbol{V}' \setminus (\boldsymbol{C} \cup \boldsymbol{Y}')$ and $V'_j \in \boldsymbol{C} \cup \boldsymbol{Y}'$ in the same district will lead to an inconsistent edge assignment. Thus the conditions of Theorem 2 imply identification of $p'(Y'(\boldsymbol{f}'_{\alpha'}))$, with the identifying formula Equation 1 following from Theorem 60 in Richardson et al. (2017).

We now show the other direction. If this first condition is violated, we are unidentified by Theorem 5.2 in Shpitser & Tchetgen (2016). If the second condition is violated, then we can take advantage of the fact that shift interventions on the treated (SITs) are a special case of the cross world policies defined here, and use the construction used in the proof of Theorem 4 in Sani et al. (2020) to show non-identification for SITs when this condition is violated. If the third condition is violated then by Theorem 7 of Shpitser & Sherman (2018), the edge policy intervention is not identified. Thus, violation of any of the conditions of Theorem 2 implies non identification. □

References

Daniel Malinsky, Ilya Shpitser, and Thomas Richardson. A potential outcomes calculus for identifying conditional path-specific effects. *AISTATS 2019 - 22nd International Conference on Artificial Intelligence and Statistics*, 89, 2020.

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

Numair Sani, Jaron J R Lee, and Ilya Shpitser. Identification and Estimation of Causal Effects Defined by Shift Interventions. *UAI*, 2020.

Ilya Shpitser and Eli Sherman. Identification of personalized effects associated with causal pathways. *34th Conference on Uncertainty in Artificial Intelligence 2018, UAI 2018*, 1:530–539, 2018. ISSN 1525-3384.

Ilya Shpitser and Eric Tchetgen Tchetgen. Causal inference with a graphical hierarchy of interventions. *Annals of statistics*, 44(6):2433, 2016.