A Proofs

Proof of Proposition 1 Under these conditions, $V(\cdot)$ is never evaluated in the recursive evaluation of $\mathbf{Y}(\mathbf{Z}=\mathbf{z}, \mathbf{A}=\mathbf{a})$ by equation (1) for any $V \in \{\mathbf{Z} \setminus \mathbf{A}\}$. $\square$

Proof of Proposition 2 By generalized consistency, $\mathbf{A}(\mathbf{z})=\mathbf{a}$ implies $\mathbf{Y}(\mathbf{z}, \mathbf{a})=\mathbf{Y}(\mathbf{z})$, and by causal irrelevance $\mathbf{Y}(\mathbf{z}, \mathbf{a})=\mathbf{Y}(\mathbf{a})$. $\square$

Proof of Proposition 3 We will show that conditions $(i),(ii),(iii)$ require that $Z(\mathbf{a})=Z(\mathbf{b})$ for all $Z \in \mathbf{X} \cup \mathbf{Y}$. It follows that if there exists $Z \in \mathbf{X} \cup \mathbf{Y}$ such that $Z(\mathbf{a}) \neq Z(\mathbf{b})$, there is no single value of $\epsilon_{\mathbf{V}}$ that leads to $\mathbf{X}(\mathbf{a})=\mathbf{x}$ and to $\mathbf{Y}(\mathbf{b})=\mathbf{y}$, and the events must be contradictory.

Let $\mathbf{C}_1$ be all variables that are causally relevant to Z in both $\mathbf{X} \cup \mathbf{A}$ and $\mathbf{Y} \cup \mathbf{B}$, let $\mathbf{C}_2$ be all variables that are causally relevant to Z in $\{\mathbf{X} \cup \mathbf{A}\} \setminus \{\mathbf{Y} \cup \mathbf{B}\}$, and that are causally relevant to Z given $\mathbf{Y} \cup \mathbf{B}$, and let $\mathbf{C}_3$ be all variables that are causally relevant to Z in $\{\mathbf{Y} \cup \mathbf{B}\} \setminus \{\mathbf{X} \cup \mathbf{A}\}$ and that are causally relevant to Z given $\mathbf{X} \cup \mathbf{A}$.

We note that condition (i) specifies that $\mathbf{C}_1(\mathbf{a})=\mathbf{C}_1(\mathbf{b})$. Then, condition (ii) requires that $\mathbf{C}_2(\mathbf{a})=\mathbf{C}_2(\mathbf{b})$; otherwise there would be a contradiction between $\mathbf{X}(\mathbf{a})=\mathbf{x}$ and $\mathbf{Y}(\mathbf{b})=\mathbf{y} \wedge \mathbf{C}_2(\mathbf{b})=\mathbf{c}_2$. In other words, there are no values of $\epsilon_{\mathbf{V}}$ that lead to $\mathbf{X}(\mathbf{a})=\mathbf{x}$ that do not lead to $\mathbf{Y}(\mathbf{b})=\mathbf{y} \wedge \mathbf{C}_2(\mathbf{b})=\mathbf{C}_2(\mathbf{a})$. For an analogous reason, condition (iii) requires that $\mathbf{C}_3(\mathbf{a})=\mathbf{C}_3(\mathbf{b})$,

We next note that by construction, no variable $D \neq Z$ in $\{\mathbf{X} \cup \mathbf{Y} \cup \mathbf{A} \cup \mathbf{B}\} \setminus \{\mathbf{C}_1 \cup \mathbf{C}_2 \cup \mathbf{C}_3\}$ is causally relevant to Z given $\mathbf{C}_1 \cup \mathbf{C}_2 \cup \mathbf{C}_3$.

Under conditions $(i),(ii),(iii)$, by consistency $Z(\mathbf{a})=Z(\mathbf{c}_1, \mathbf{c}_2, \mathbf{c}_3, \mathbf{a}, \mathbf{x} \backslash\{z\})$. By causal irrelevance of all variables $D \neq Z$ in $\mathbf{A} \cup \mathbf{X}$ not in $\mathbf{C}_1 \cup \mathbf{C}_2 \cup \mathbf{C}_3$ given $\mathbf{C}_1 \cup \mathbf{C}_2 \cup \mathbf{C}_3$ we have $Z(\mathbf{c}_1, \mathbf{c}_2, \mathbf{c}_3, \mathbf{a}, \mathbf{x} \backslash\{z\})=Z(\mathbf{c}_1, \mathbf{c}_2, \mathbf{c}_3)$. For the same reasons, $Z(\mathbf{b})=Z(\mathbf{c}_1, \mathbf{c}_2, \mathbf{c}_3, \mathbf{b}, \mathbf{y} \backslash\{z\})=Z(\mathbf{c}_1, \mathbf{c}_2, \mathbf{c}_3)$. This yields $Z(\mathbf{a})=Z(\mathbf{b})$, completing the proof. This recursive definition of contradiction between events must resolve because each recursive call expands the number of variables specified in one of the events, and there are a finite number of variables in the graph. $\square$

Proof of Proposition 5 Each event in the partition either directly specifies $\mathbf{Y}(\mathbf{a_1})=\mathbf{y}$, or specifies an event that implies $\mathbf{Y}(\mathbf{a_1})=\mathbf{y}$ by Proposition 2, so we know the disjunction of these events is a subset of $\mathbf{Y}(\mathbf{a_1})=\mathbf{y}$.

Now we show that all sets are disjoint. Event (13) and all events of the form of (14) are pairwise disjoint, as each requires a different event under the intervention $\mathbf{Z}=\mathbf{z_1}$. These events are all disjoint from event (12), as the former each specifies a value of $\mathbf{z}'$ for which $\mathbf{A}(\mathbf{z}')=\mathbf{a}$, and the latter specifies that no such $\mathbf{z}'$ can exist.

Finally, we show that the events are exhaustive. In addition to the requirement that $\mathbf{Y}(\mathbf{a_1})=\mathbf{y}$ specified above, a disjunction of the partition set events requires that $\mathbf{A}(\mathbf{z_1})$ take a value in $\mathbf{a_1}, \cdots, \mathbf{a_N}$, which is tautological. It also requires that $\forall \mathbf{z}'\big(\mathbf{A}(\mathbf{z}') \neq \mathbf{a_1}\big)$ or $\exists \mathbf{z}'\big(\mathbf{A}(\mathbf{z}')=\mathbf{a_1}\big)$, which is likewise tautological. No other requirements are present. $\square$

Proof of Proposition 6 $E_{\mathbf{x}} \wedge Y(b) \neq y \implies \psi_{\mathbf{b}}(E_{\mathbf{x}}) \wedge Y(\mathbf{b}) \neq \mathbf{y}$. Therefore $P(E_{\mathbf{x}}, Y(b) \neq y) \leq P(\psi_{\mathbf{b}}(E_{\mathbf{x}}), Y(\mathbf{b}) \neq \mathbf{y})$.

$P(E_{\mathbf{x}})=P(E_{\mathbf{x}}, \mathbf{Y}(\mathbf{b})=\mathbf{y})+P(E_{\mathbf{x}}, \mathbf{Y}(\mathbf{b}) \neq \mathbf{y})$, which, in combination with the preceding, yields $P(E_{\mathbf{x}}, \mathbf{Y}(\mathbf{b})=\mathbf{y}) \geq P(E_{\mathbf{x}})-P(\psi_{\mathbf{b}}(E_{\mathbf{x}}), \mathbf{Y}(\mathbf{b}) \neq \mathbf{y})$.

By construction, $E_{\mathbf{x}} \implies \mathbf{X}(\mathbf{a})=\mathbf{x}$, yielding the first bound. A symmetric argument yields the second. $\square$

Proof of Proposition 7 The inclusion-exclusion formula states that $\sum_{i=1}^N P(E_i)=P(\cup_{i=1}^N E_i)+\sum_{i<j} P(E_i \cap E_j)$. Under the conditions of the Proposition, $\sum_{i<j} P(E_i \cap E_j) \leq k-1$, as any set of values of $\epsilon_{\mathbf{V}}$ may imply at most $k-1$ of the events in $\{E_i \cap E_j \mid i<j\}$, and the set of all possible values of $\epsilon_{\mathbf{V}}$ has measure 1. Noting that $P(\cup_{i=1}^N E_i) \leq 1$ by definition, we have $P(\cup_{i=1}^N E_i)+\sum_{i<j} P(E_i \cap E_j) \leq k$. $\square$

Proof of Corollary 2 $\tilde{\mathbf{A}}$ is causally irrelevant to $\mathbf{Y}$ given $\mathbf{A}$, so the lower bounds follow directly from Propositions 2 and 4. The upper bounds follow because, by the same reasoning, $P(\mathbf{Y}(\tilde{\mathbf{a}}) \neq \mathbf{y}, \hat{\mathbf{A}}(\tilde{\mathbf{a}})=\hat{\mathbf{a}}) \leq P(\mathbf{Y}(\mathbf{a}) \neq \mathbf{y})$. $\square$

Proof of Corollary 4 Proposition 3 tells us that if $\mathbf{Z}$ is a generalized instrument for $\mathbf{A}$ with respect to $\mathbf{Y}$, two events of the form $\mathbf{A}(\mathbf{z})=\mathbf{a} \wedge \mathbf{Y}(\mathbf{a})=\mathbf{y}$ and $\mathbf{A}(\mathbf{z}')=\mathbf{a}' \wedge \mathbf{Y}(\mathbf{z}')=\mathbf{y}'$ are contradictory if $\neg\big((\mathbf{z}=\mathbf{z}' \wedge \mathbf{a} \neq \mathbf{a}') \vee (\mathbf{a}=\mathbf{a}' \wedge \mathbf{y} \neq \mathbf{y}')\big)$. Therefore $\Phi(S)$ provides the size of the largest subset of events that are mutually compatible. The result then follows immediately from Proposition 7. $\square$

B Equivalence Class Completeness

In this appendix we introduce an assumption we call *equivalence class completeness*. Following [1], we say two values in the domain of $\epsilon_{\mathbf{V}}$ are in the same equivalence class if they will produce the same results through equation (1) for every variable under every intervention.

Assumption 1 (Equivalence Class Completeness)**.** *Every equivalence class of $\epsilon_{\mathbf{V}}$ is non-empty in the domain of $\epsilon_{\mathbf{V}}$.*

We note that this assumption precludes the possibility of vacuous edges (edges that reflect no causal influence) in the graph. In the case of a vacuous edge $X \rightarrow V$, for example, there will be no values of ϵ_V that produce $V = v$ for some setting of $X = x$ and $V = v'$ for another setting of $X = x'$, because V is not a function of X. This would mean that the equivalence class of values of $\epsilon_{\mathbf{V}}$ that lead to $V = v$ under intervention $X = x$ and $V = v'$ under intervention $X = x'$ is empty, violating the assumption.

For the same reason, this assumption precludes the possibility of context-specific exclusion restrictions. If there is an edge $X \rightarrow V$ such that for some level of the other parents of V, denoted by $\mathbf{Y} = \mathbf{y}$, V is not a function of X, then there will exist no value of ϵ_V that leads to $V = v$ under intervention $X = x, \mathbf{Y} = \mathbf{y}$ but to $V = v'$ under intervention $X = x', \mathbf{Y} = \mathbf{y}$.

We now show that under this assumption, the criteria in Propositions 2 and 3 for cross-world implication and event contradiction respectively are necessary as well as sufficient. It follows that unless there exists background knowledge that the equivalence class completeness assumption is violated, due for example to deterministic causal relationships, all implications and contradictions relevant for deriving bounds and inequality constraints can be obtained using these criteria.

Proposition 8. *Under the equivalence class completeness assumption, $\mathbf{Z}$ is causally irrelevant to $\mathbf{Y}$ given $\mathbf{A}$ if and only if:*

$$\mathbf{A}(\mathbf{z}) = \mathbf{a} \wedge \mathbf{Y}(\mathbf{z}) = \mathbf{y} \implies \mathbf{Y}(\mathbf{a}) = \mathbf{y}.$$

Proof. Sufficiency is given by proposition 2. We demonstrate necessity as follows. Assume $\mathbf{Z}$ is not causally irrelevant to $\mathbf{Y}$ given $\mathbf{A}$, i.e. there is a path from $\mathbf{Z}$ to $\mathbf{Y}$ not through $\mathbf{A}$. Then by equivalence class completeness, there must be values of $\epsilon_{\mathbf{V}}$ for which $\mathbf{Y}$ is a function of $\mathbf{Z}$ when $\mathbf{A}$ is exogenously set and $\mathbf{Z}$ does not take the value $\mathbf{z}$ under no intervention. Therefore, there will exist values of $\epsilon_{\mathbf{V}}$ such that $\mathbf{Y}(\mathbf{a}, \mathbf{z}) \neq \mathbf{Y}(\mathbf{a})$. By generalized consistency $\mathbf{A}(\mathbf{z}) = \mathbf{a} \wedge \mathbf{Y}(\mathbf{z}) = \mathbf{y} \implies \mathbf{Y}(\mathbf{a}, \mathbf{z}) = \mathbf{y}$, which contradicts $\mathbf{A}(\mathbf{z}) = \mathbf{a} \wedge \mathbf{Y}(\mathbf{z}) = \mathbf{y} \implies \mathbf{Y}(\mathbf{a}) = \mathbf{y}$. ☐

Proposition 9. *Under the equivalence class completeness assumption, two events $\mathbf{X}(\mathbf{a}) = \mathbf{x}$ and $\mathbf{Y}(\mathbf{b}) = \mathbf{y}$ are contradictory if and only if there exists $Z \in \mathbf{X} \cup \mathbf{Y}$ such that $Z(\mathbf{a}) \neq Z(\mathbf{b})$, and all of the following hold:*

(i) Variables in the subsets of both $\mathbf{X} \cup \mathbf{A}$ and $\mathbf{Y} \cup \mathbf{B}$ causally relevant for Z are set to the same values in $\mathbf{x}, \mathbf{a}$, and $\mathbf{y}, \mathbf{b}$.

(ii) Let $C \in \{\mathbf{X} \cup \mathbf{A}\} \setminus \{\mathbf{Y} \cup \mathbf{B}\}$ be any variable that is causally relevant to Z in $\mathbf{X} \cup \mathbf{A}$ and causally relevant to Z given $\mathbf{Y} \cup \mathbf{B}$, with C set to c in $\mathbf{x}, \mathbf{a}$. Then $\mathbf{X}(\mathbf{a}) = \mathbf{x}$ and $\mathbf{Y}(\mathbf{b}) = \mathbf{y} \wedge C(\mathbf{b}) = c'$ are contradictory when $c \neq c'$.

(iii) Let $C \in \{\mathbf{Y} \cup \mathbf{B}\} \setminus \{\mathbf{X} \cup \mathbf{A}\}$ be any variable that is causally relevant to Z in $\mathbf{Y} \cup \mathbf{B}$ and causally relevant to Z given $\mathbf{X} \cup \mathbf{A}$, with C set to c in $\mathbf{y}, \mathbf{b}$. Then $\mathbf{Y}(\mathbf{b}) = \mathbf{y}$ and $\mathbf{X}(\mathbf{a}) = \mathbf{x} \wedge C(\mathbf{a}) = c'$ are contradictory when $c \neq c'$.

Proof. Sufficiency is given by Proposition 3. To see the necessity of condition (i), we note that if variables causally relevant in $\mathbf{X} \cup \mathbf{A}$ and in $\mathbf{Y} \cup \mathbf{B}$ took different values in $\mathbf{x}, \mathbf{a}$ and $\mathbf{y}, \mathbf{b}$, then if $Z(\mathbf{x}, \mathbf{a}) \neq Z(\mathbf{y}, \mathbf{b})$, there must be an equivalence class that leads to these two results under their respective interventions. By the equivalence class completeness assumption it will be non-empty. Therefore there exists a value of $\epsilon_{\mathbf{V}}$ that leads to both events, and they are not contradictory.

We now demonstrate the necessity of condition (ii). If (ii) does not hold, there must be a variable D that is causally relevant to Z in $\mathbf{X} \cup \mathbf{A}$ and given $\mathbf{Y} \cup \mathbf{B}$ that can take different values under equivalence classes of $\epsilon_{\mathbf{V}}$ that lead to $\mathbf{X}(\mathbf{a}) = \mathbf{a}$ and $\mathbf{Y}(\mathbf{b}) = \mathbf{y}$ under their respective interventions. Because D is causally relevant given both the remainder of $\mathbf{X} \cup \mathbf{A}$, and given all of $\mathbf{Y} \cup \mathbf{B}$, and can for single value of $\epsilon_{\mathbf{V}}$ take different values under the relevant interventions, it is possible for that value of $\epsilon_{\mathbf{V}}$ to yield different values of Z under the two interventions. By equivalence class completeness, an $\epsilon_{\mathbf{V}}$ leading to this result must exist, leading to a lack of contradiction between the two events. Condition (iii) is necessary by an analogous argument. ☐

C Redundant Lower Bounds

We present results that establish the redundance of lower bounds induced by certain events E_1 and E_2 through Corollary 3.

We first observe that the event chosen for E_1 in Proposition 6 should be compatible with the event $\mathbf{Y}(\mathbf{b}) = \mathbf{y}$. If it is not, $\psi_{\mathbf{b}}(E_1) \wedge \mathbf{Y}(\mathbf{b}) \neq \mathbf{y}$ is equivalent to $\psi_{\mathbf{b}}(E_1)$.

Because $E_1 \implies \psi_{\mathbf{b}}(E_1)$, by Proposition 4 any such E_1 will induce a negative lower bound, which is of course uninformative. An analogous argument can be made for E_2.

We next consider a proposition that explains why we did not need to consider the bound induced by $E_1 \triangleq A(\bar{z}) = \bar{a}$ to obtain sharp bounds in Section 3.

Proposition 10. *Let E_1 imply $\mathbf{X}(\mathbf{a}) = \mathbf{a}$ and let $\mathbf{Y}(\mathbf{b}) \neq \mathbf{y}$ imply $\psi_{\mathbf{b}}(E_1)$. Then the event $E_2 \triangleq \mathbf{Y}(\mathbf{b}) = \mathbf{y}$ induces, through Proposition 6, a weakly better bound than does E_1. An analogous claim holds for E_2.*

Proof. $P(\psi_{\mathbf{b}}(E_1), \mathbf{Y}(\mathbf{b}) \neq \mathbf{y}) = P(\mathbf{Y}(\mathbf{b}) \neq \mathbf{y})$, as $\mathbf{Y}(\mathbf{b}) \neq \mathbf{y} \implies \psi_{\mathbf{b}}(E_1)$ by assumption.

The lower bound induced by E_1, given by Proposition 6 as $P(E_1) - P(\psi_{\mathbf{b}}(E_1), \mathbf{Y}(\mathbf{b}) \neq \mathbf{b})$, can now be expressed as $P(\mathbf{Y}(\mathbf{b}) = \mathbf{y}) - P(\neg E_1)$.

We now note $\psi_{\mathbf{a}}(\mathbf{Y}(\mathbf{b}) = \mathbf{b})) \wedge \mathbf{X}(\mathbf{a}) \neq \mathbf{x} \implies \neg E_1$, as $E_1 \implies \mathbf{X}(\mathbf{a}) = \mathbf{x}$ by construction. Therefore $P(\neg E_1) \geq P(\psi_{\mathbf{a}}(\mathbf{Y}(\mathbf{b}) = \mathbf{b}) \wedge \mathbf{X}(\mathbf{a}) \neq \mathbf{x})$, and the bound induced by $E_2 \triangleq \mathbf{Y}(\mathbf{b}) = \mathbf{y}$, given by Proposition 6 as $P(\mathbf{Y}(\mathbf{b}) = \mathbf{y}) - P(\psi_{\mathbf{a}}(\mathbf{Y}(\mathbf{b}) = \mathbf{b})) \wedge \mathbf{X}(\mathbf{a}) \neq \mathbf{x})$, must be better than that induced by E_1. ☐

In the binary IV case described in Section 3, every event under intervention $Z = z$ is compatible with the event $A(\bar{z}) = \bar{a}$. This means that in particular $\psi_z(A(\bar{z}) = \bar{a})$ is implied by $\neg\big(A(z) = a \wedge Y(z) = y\big)$. The lower bound on event (4) induced by $E_1 \triangleq A(\bar{z}) = \bar{a}$ is therefore redundant given the bound induced by $E_2 \triangleq A(z) = a \wedge Y(z) = y$.

Next, we identify an additional condition under which bounds induced by particular valid choices of E_1 and E_2 are irrelevant. This condition does not appear in the IV model.

Proposition 11. *If two candidates for events E_1 (E_2), under Proposition 6 are each compatible with the same events under $\mathbf{B} = \mathbf{b}$ ($\mathbf{A} = \mathbf{a}$), the candidate event with larger density will induce a better bound.*

Proof. The bound in Proposition 6 is expressed as the density of E_1 (E_2) less a function of the events compatible with E_1 (E_2). If the two candidate events are compatible with the same set of events, the negative quantity in the bound will be the same. The bound with the larger positive quantity – the density of E_1 (E_2) – must be larger. ☐

Proposition 12. *Under the equivalence class completeness assumption, an event $\mathbf{Y}(\mathbf{a}) = \mathbf{y} \wedge \mathbf{X}(\mathbf{a}) = \mathbf{x}$ is compatible with the same events under intervention $\mathbf{A} = \mathbf{a}'$ as is $\mathbf{X}(\mathbf{a}) = \mathbf{x}$ if and only if $\mathbf{Y}$, and all descendants of $\mathbf{Y}$ in $\mathbf{X}$ to which $\mathbf{Y}$ is causally relevant given the remainder of $\mathbf{X}$, have at least one causally relevant ancestor in $\mathbf{A}$ that takes different values in $\mathbf{a}$ than in $\mathbf{a}'$.*

Proof. If an event does not contradict $\mathbf{Y}(\mathbf{a}) = \mathbf{y} \wedge \mathbf{X}(\mathbf{a}) = \mathbf{x}$, it will not contradict the less restrictive event $\mathbf{X}(\mathbf{a}) = \mathbf{x}$.

We consider an event compatible with $\mathbf{X}(\mathbf{a}) = \mathbf{x}$. By Proposition 9, if it is to contradict $\mathbf{X}(\mathbf{a}) = \mathbf{x} \wedge \mathbf{Y}(\mathbf{a}) = \mathbf{y}$ under the equivalence class completeness assumption, then there must be a variable Z satisfying the conditions of that proposition. This Z cannot be in $\mathbf{Y}$, or any of its descendants in $\mathbf{X}$ to which it is causally relevant given the remainder of $\mathbf{X}$, by the condition that they each have a causally relevant ancestor in $\mathbf{A}$ that differs between $\mathbf{a}$ and $\mathbf{a}'$. If it is any variable to which $\mathbf{Y}$ is not causally relevant, then the causally relevant ancestors are the same in $\mathbf{X}(\mathbf{a}) = \mathbf{x} \wedge \mathbf{Y}(\mathbf{a}) = \mathbf{y}$ as in $\mathbf{X}(\mathbf{a}) = \mathbf{x}$, so the event must also be compatible with $\mathbf{X}(\mathbf{a}) = \mathbf{x} \wedge \mathbf{Y}(\mathbf{a}) = \mathbf{y}$ if it is compatible with $\mathbf{X}(\mathbf{a}) = \mathbf{x}$.

Finally, we demonstrate the necessity of these conditions. If they failed to hold, some variable in $\mathbf{Y} \cup \mathbf{X}$ in $\mathbf{Y}$ or to which $\mathbf{Y}$ is causally relevant given the remainder of $\mathbf{X}$ would have no causally relevant ancestor in $\mathbf{A}$ that differed under the two interventions. We call such a variable Z, and say it takes value z. Then we construct the event $\mathbf{X}'(\mathbf{a}') = \mathbf{x}' \wedge \mathbf{Y}'(\mathbf{a}') = \mathbf{y}' \wedge \mathbf{Z}(\mathbf{a}') \neq \mathbf{z}$, with $\mathbf{X}', \mathbf{Y}'$ denoting $\mathbf{X} \setminus \{Z\}, \mathbf{Y} \setminus \{Z\}$. This event contradicts $\mathbf{Y}(\mathbf{a}) = \mathbf{y} \wedge \mathbf{X}(\mathbf{a}) = \mathbf{x}$ but does not contradict $\mathbf{X}(\mathbf{a}) = \mathbf{x}$ by Proposition 9. ☐

We note that because the bounds derived by Corollary 3 do not make use of any additional implications that may result from violations of the equivalence class completeness assumption, these results lead directly to the following Corollary:

Corollary 5. *Let the event $\tilde{E} \triangleq \mathbf{Y}(\mathbf{a}) = \mathbf{y} \wedge \mathbf{X}(\mathbf{a}) = \mathbf{x}$ be compatible with the same events under $\mathbf{A} = \mathbf{a}'$ as $\mathbf{X}(\mathbf{a}) = \mathbf{x}$ by Proposition 12, and be a valid candidate for E_1 (E_2). Then by Proposition 11, any event $\mathbf{X}(\mathbf{a}) = \mathbf{x} \wedge \mathbf{W}(\mathbf{a}) = \mathbf{y}_{\mathbf{w}}$, with $\mathbf{W} \subset \mathbf{Y}$, that is also a valid candidate for E_1 (E_2) will induce a better bound through Corollary 3 than $\tilde{E}$.*

D Numerical Examples of Bound Width

In this section, we provide numerical examples of bounds in two models. These examples demonstrate that bounds tend to be most informative when the instrument and treatment are highly correlated. It is our hope that they

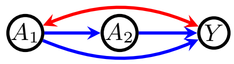

Figure 4: The Inclusive Frontdoor Model

will provide some intuition about when these bounds will be of use.

Consider the causal model described by the graph in Fig. 4. Suppose all variables are binary and we are interested in the probability $P(Y(A_1 = 1, A_2 = 1) = 1)$. If we observe the following probabilities,

$$P(A_1 = 1, A_2 = 1, Y = 1) = .01$$
$$P(A_1 = 1, A_2 = 1, Y = 0) = .08,$$

we can use Corollary 1 to obtain the bounds

$$.01 \leq P(Y(A_1 = 1, A_2 = 1) = 1) \leq .92.$$

These bounds are quite wide, and unlikely to be informative. Now suppose that we observe the following conditional probability

$$P(A_2 = 1 \mid A_1 = 1) = .1. \quad (23)$$

Noting that $P(A_1(a_2) = a_1, Y(a_2) = y)$ is identified as $\frac{P(a_1,a_2,y)}{P(a_2|a_1)}$, we can now use Corollary 2 to obtain the bounds

$$.1 \leq P(Y(A_1 = 1, A_2 = 1) = 1) \leq .2.$$

These bounds are much tighter, and exclude .5, which may be important in some cases. If instead we observe the conditional probability

$$P(A_2 = 1 \mid A_1 = 1) = .5, \quad (24)$$

then Corollary 2 yields the much less informative bounds

$$.02 \leq P(Y(A_1 = 1, A_2 = 1) = 1) \leq .84.$$

In this example A_2 is used as a generalized instrument for A_1, as discussed in Section 4. The tightness of the bounds therefore depends on the relationship between the two, as demonstrated by the differences in bounds under the conditional probabilities (23) and (24).

We now consider the IV Model with Covariates, depicted in Fig. 3 (a), and discussed in Section 6. To build an understanding of the utility of our bounds, we randomly generated distributions from the model. These distributions were generated by sampling the parameter for each observed binary random variable, conditional on each setting of its parents, from a symmetric Beta distribution, with parameters equal to 1. The unobserved variable U was assumed to have cardinality 16, to allow for every possible equivalence class [1], and its distribution was drawn

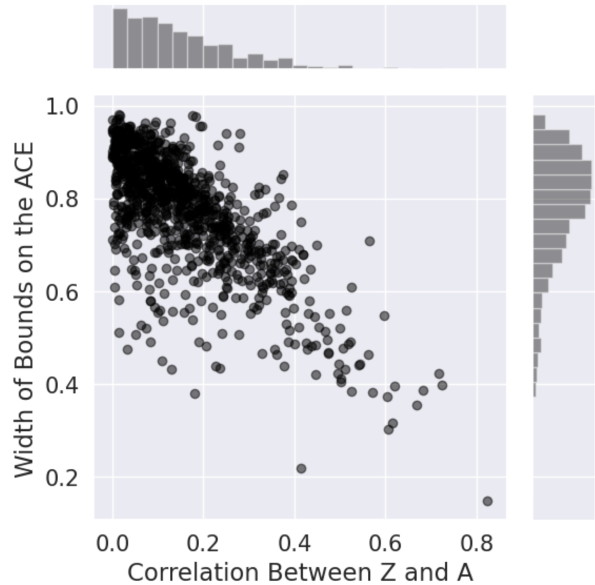

Figure 5: Each point represents a randomly generated distribution in the IV model with covariates, depicted in Fig. 3 (a). This plot shows informally that correlation between the generalized instrument Z and the treatment A is associated with tighter bounds on the ACE, $E[Y(A = 1) - Y(A = 0)]$. The histograms show the marginal distributions of correlation and bound width.

from a symmetric Dirichlet distribution with parameters equal to 0.1.

We then calculated the correlation between A and Z, as well as bounds on the ACE, $E[Y(A = 1) - Y(A = 0)]$, for each distribution. The results are presented in Fig. 5.

We observe that, as expected, greater correlation between the generalized instrument and the treatment is associated with tighter bounds. This pattern persisted across simulations for additional models, and across various approaches to sampling distributions from the model.

The marginal distributions of correlation and bound width, shown as histograms in Fig. 5 seem to be sensitive to the approach used to sample distributions from the model. For distributions sampled as described above, and used to generate Fig. 5, the mean width of bounds on the ACE was 0.77, with a standard deviation of 0.12. In 4% of the distributions, the bounds excluded 0. We find that these values are also sensitive to changes in how distributions were sampled.

E Bounds on $P(Y(\bar{a}) = \bar{y})$ in the IV Model With Covariates

The remainder of this section is a LaTeX friendly printout of the steps taken by our implementation of the algorithm described in this work when applied to bounding $P(\bar{y}(\bar{a}))$ in the IV model with covariates, as described in Section 6.

To begin, we partition $\bar{y}(\bar{a})$ as described in Proposition 5:

$$\bar{y}(\bar{a}) \wedge a(\bar{z}) \wedge a(z)$$
$$\bar{a}(\bar{z}) \wedge \bar{y}(\bar{z})$$
$$(a(\bar{z}) \wedge (\bar{a}(z) \wedge \bar{y}(z))).$$

As before, no lower bound is provided for the first event in the partition, and the second has an identified density. We now consider the last event in the partition $a(\bar{z}) \wedge (\bar{a}(z) \wedge \bar{y}(z))$.

The following events imply $a(\bar{z})$ and can therefore be used as E_1 events in Corollary 3:

$$a(\bar{z}) \wedge \bar{y}(\bar{z})$$
$$a(\bar{z}) \wedge y(\bar{z})$$
$$a(\bar{z}) \wedge \bar{c}(\bar{z})$$
$$a(\bar{z}) \wedge c(\bar{z})$$
$$a(\bar{z}) \wedge \bar{y}(\bar{z}) \wedge \bar{c}(\bar{z})$$
$$a(\bar{z}) \wedge \bar{y}(\bar{z}) \wedge c(\bar{z})$$
$$a(\bar{z}) \wedge y(\bar{z}) \wedge \bar{c}(\bar{z})$$
$$a(\bar{z}) \wedge y(\bar{z}) \wedge c(\bar{z}).$$

Likewise, the following events imply $\bar{a}(z) \wedge \bar{y}(z)$ and can therefore be used as E_2 events in Corollary 3:

$$\bar{a}(z) \wedge \bar{y}(z)$$
$$\bar{a}(z) \wedge \bar{y}(z) \wedge \bar{c}(z)$$
$$\bar{a}(z) \wedge \bar{y}(z) \wedge c(z)$$

By Proposition 10, $a(\bar{z})$, which implies $a(\bar{z})$, and therefore would be a candidate for use as an E_1 event, is redundant.

We now iterate through each potential event for $E_1$and E_2, examining the resulting bound.

The event $a(\bar{z}) \wedge \bar{y}(\bar{z})$ is compatible with $\bar{y}(z) \vee (\bar{a}(z) \wedge y(z))$. Therefore to compute the bound induced by using it as an E_1 event, we must subtract from its density the portion of this compatible event that does not entail the negation of $\bar{a}(z) \wedge \bar{y}(z)$. This portion is $(a(z) \wedge \bar{y}(z)) \vee (\bar{a}(z) \wedge y(z))$, yielding the bound $P_{\bar{z}}(a, \bar{y}) - \big(P_z(a, \bar{y}) + P_z(\bar{a}, y)\big)$.

The event $a(\bar{z}) \wedge y(\bar{z})$ is compatible with $(\bar{a}(z) \wedge \bar{y}(z)) \vee y(z)$. Therefore to compute the bound induced by using it as an E_1 event, we must subtract from its density the portion of this compatible event that does not entail the negation of $\bar{a}(z) \wedge \bar{y}(z)$. This portion is $y(z)$, yielding the bound $P_{\bar{z}}(a, y) - P_z(y)$.

The event $a(\bar{z}) \wedge \bar{c}(\bar{z})$ is compatible with $(a(z) \wedge \bar{c}(z)) \vee c(z)$. Therefore to compute the bound induced by using it as an E_1 event, we must subtract from its density the portion of this compatible event that does not entail the negation of $\bar{a}(z) \wedge \bar{y}(z)$. This portion is $(a(z) \wedge y(z) \wedge \bar{c}(z)) \vee (a(z) \wedge \bar{y}(z)) \vee (y(z) \wedge c(z))$, yielding the bound $P_{\bar{z}}(a, \bar{c}) - \big(P_z(a, y, \bar{c}) + P_z(a, \bar{y}) + P_z(y, c)\big)$.

The event $a(\bar{z}) \wedge c(\bar{z})$ is compatible with $\bar{c}(z) \vee (a(z) \wedge c(z))$. Therefore to compute the bound induced by using it as an E_1 event, we must subtract from its density the portion of this compatible event that does not entail the negation of $\bar{a}(z) \wedge \bar{y}(z)$. This portion is $(y(z) \wedge \bar{c}(z)) \vee (a(z) \wedge \bar{y}(z) \wedge \bar{c}(z)) \vee (a(z) \wedge c(z))$, yielding the bound $P_{\bar{z}}(a, c) - \big(P_z(y, \bar{c}) + P_z(a, \bar{y}, \bar{c}) + P_z(a, c)\big)$.

The event $a(\bar{z}) \wedge \bar{y}(\bar{z}) \wedge \bar{c}(\bar{z})$ is compatible with $(\bar{a}(z) \wedge y(z) \wedge c(z)) \vee (a(z) \wedge \bar{y}(z) \wedge \bar{c}(z)) \vee (\bar{y}(z) \wedge c(z))$. Therefore to compute the bound induced by using it as an E_1 event, we must subtract from its density the portion of this compatible event that does not entail the negation of $\bar{a}(z) \wedge \bar{y}(z)$. This portion is $(\bar{a}(z) \wedge y(z) \wedge c(z)) \vee (a(z) \wedge \bar{y}(z))$, yielding the bound $P_{\bar{z}}(a, \bar{y}, \bar{c}) - \big(P_z(\bar{a}, y, c) + P_z(a, \bar{y})\big)$.

The event $a(\bar{z}) \wedge \bar{y}(\bar{z}) \wedge c(\bar{z})$ is compatible with $(\bar{y}(z) \wedge \bar{c}(z)) \vee (\bar{a}(z) \wedge y(z) \wedge \bar{c}(z)) \vee (a(z) \wedge \bar{y}(z) \wedge c(z))$. Therefore to compute the bound induced by using it as an E_1 event, we must subtract from its density the portion of this compatible event that does not entail the negation of $\bar{a}(z) \wedge \bar{y}(z)$. This portion is $(\bar{a}(z) \wedge y(z) \wedge \bar{c}(z)) \vee (a(z) \wedge \bar{y}(z))$, yielding the bound $P_{\bar{z}}(a, \bar{y}, c) - \big(P_z(\bar{a}, y, \bar{c}) + P_z(a, \bar{y})\big)$.

The event $a(\bar{z}) \wedge y(\bar{z}) \wedge \bar{c}(\bar{z})$ is compatible with $(a(z) \wedge y(z) \wedge \bar{c}(z)) \vee (\bar{a}(z) \wedge \bar{y}(z) \wedge c(z)) \vee (y(z) \wedge c(z))$. Therefore to compute the bound induced by using it as an E_1 event, we must subtract from its density the portion of this compatible event that does not entail the negation of $\bar{a}(z) \wedge \bar{y}(z)$. This portion is $(a(z) \wedge y(z) \wedge \bar{c}(z)) \vee (y(z) \wedge c(z))$, yielding the bound $P_{\bar{z}}(a, y, \bar{c}) - \big(P_z(a, y, \bar{c}) + P_z(y, c)\big)$.

The event $a(\bar{z}) \wedge y(\bar{z}) \wedge c(\bar{z})$ is compatible with $(y(z) \wedge \bar{c}(z)) \vee (\bar{a}(z) \wedge \bar{y}(z) \wedge \bar{c}(z)) \vee (a(z) \wedge y(z) \wedge c(z))$. Therefore to compute the bound induced by using it as an E_1 event, we must subtract from its density the portion of this compatible event that does not entail the negation of $\bar{a}(z) \wedge \bar{y}(z)$. This portion is $(y(z) \wedge \bar{c}(z)) \vee (a(z) \wedge y(z) \wedge c(z))$, yielding the bound $P_{\bar{z}}(a, y, c) - \big(P_z(y, \bar{c}) + P_z(a, y, c)\big)$.

The event $\bar{a}(z) \wedge \bar{y}(z)$is compatible with $\bar{y}(\bar{z}) \vee (a(\bar{z}) \wedge y(\bar{z}))$. Therefore to compute the bound induced by using it as an E_2 event, we must subtract from its density the portion of this compatible event that does not entail the negation of $a(\bar{z})$. This portion is $(\bar{a}(\bar{z}) \wedge \bar{y}(\bar{z}))$, yielding the bound $P_z(\bar{a}, \bar{y}) - P_{\bar{z}}(\bar{a}, \bar{y})$.

The event $\bar{a}(z) \wedge \bar{y}(z) \wedge \bar{c}(z)$is compatible with $(\bar{y}(\bar{z}) \wedge c(\bar{z})) \vee (a(\bar{z}) \wedge y(\bar{z}) \wedge c(\bar{z})) \vee (\bar{a}(\bar{z}) \wedge \bar{y}(\bar{z}) \wedge \bar{c}(\bar{z}))$. Therefore to compute the bound induced by using it as an E_2 event, we must subtract from its density the portion of this compatible event that does not entail the negation of $a(\bar{z})$. This portion is $(\bar{a}(\bar{z}) \wedge \bar{y}(\bar{z}))$, yielding the bound $P_z(\bar{a}, \bar{y}, \bar{c}) - P_{\bar{z}}(\bar{a}, \bar{y})$.

The event $\bar{a}(z) \wedge \bar{y}(z) \wedge c(z)$is compatible with $(\bar{a}(\bar{z}) \wedge \bar{y}(\bar{z}) \wedge c(\bar{z})) \vee (a(\bar{z}) \wedge y(\bar{z}) \wedge \bar{c}(\bar{z})) \vee (\bar{y}(\bar{z}) \wedge \bar{c}(\bar{z}))$. Therefore to compute the bound induced by using it as an E_2 event, we must subtract from its density the portion of this compatible event that does not entail the negation of $a(\bar{z})$. This portion is $(\bar{a}(\bar{z}) \wedge \bar{y}(\bar{z}))$, yielding the bound $P_z(\bar{a}, \bar{y}, c) - P_{\bar{z}}(\bar{a}, \bar{y})$.

This concludes the derivation of the bounds presented for the IV model with covariates in Section 6.

References

[1] Alexander Abraham Balke. *Probabilistic Counterfactuals: Semantics, Computation, and Applications*. PhD thesis, USA, 1996.