Single World Intervention Graphs (SWIGs)

James Robins and Thomas Richardson

·Potential outcomes are extensively used within Statistics for reasoning about causation.

·Directed acyclic graphs (DAGs) are another formalism used to represent causal systems also extensively used in Computer Science, Bioinformatics, Sociology and Epidemiology.

·Given the utility of both approaches – as demonstrated by many applications – it is natural to to wish to unify them

·idea of splitting nodes.the key

·Counterfactuals $Y(x)$ rather than Y_x

·Only treatment variables must have well defined counterfactuals. Well-defined vs Vague

Most important

·Finally graphical and counterfactual people can speak the same language and understand one another

Also

· Improves reasoning about counterfactuals on graphs as unlike other systems

twin networks (Balke Pearl) and counterfactual graphs (Shpitser Pearl)

no determinism on graphs so d-connection implies dependence (for some distribution ion the model)

DAGS

- $(V_0, V_1, \ldots, V_M)$

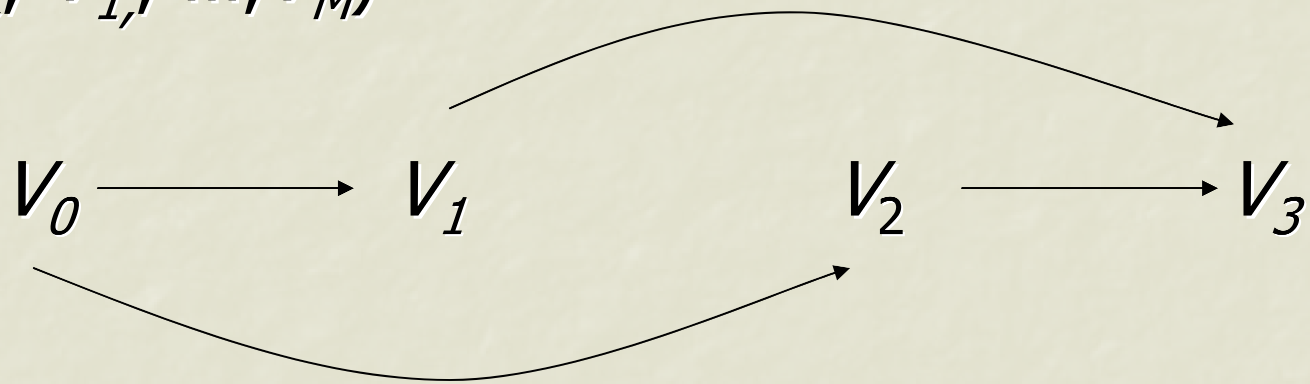

- Statistical Dag:

$$f(V) = \prod_{m=1}^{M} f(V_m | pa_m) \quad pa_m \text{ are the parents of } V_m$$

- Example

$$f(V) = f(V_3 | V_1, V_2)\, f(V_2 | V_0)\, f(V_1 | V_0)\, f(V_0)$$

DAGS

- $(V_0, V_1, \ldots, V_M)$

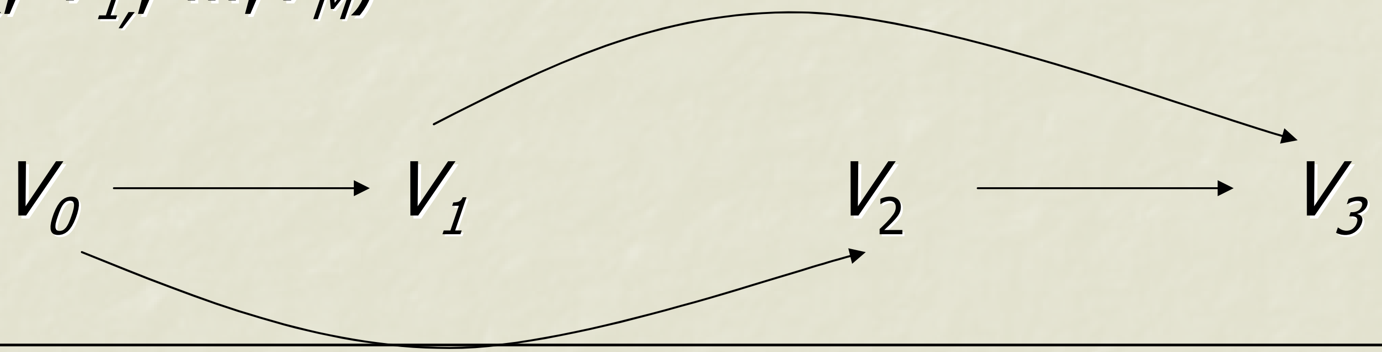

statistical dag: equivalent definitions: Each variable is independent of all variables in its past given its parents
Each variable is independent of its nondescendants given its parents.
In this DAG V2 is independent of V1 given (within each level of) V0.

- Example

$$f(V) = f(V_3|V_1, V_2)\, f(V_2|V_0)\, f(V_1|V_0)\, f(V_0)$$

V2 is independent of V1 ie OR=1
with each level of V0 if the the probabilty V2 takes any value is not predicted by V1.

Given V0, V1 is not an independent predictor of (risk factor for) V2

- Complete Dag

$$f(V) = f(V_3|V_0, V_1, V_2)\, f(V_2|V_0, V_1)\, f(V_1|V_0)\, f(V_0)$$

Dr. James M. Robins
Dr. Miguel A. Hernán

Epi. 207a — Notes on d-separation —Fall 2001

Unconditional Independence: First let's consider whether a variable A is d-separated (independent) of a variable B on a DAG unconditionally. They are independent if there is no unblocked (active) path between them down which probability can flow. To determine this, define a path as follows. Take your graph, mentally erase (for the moment only) all arrowheads. So the graph now just has edges (arcs) between nodes (variables). A path between A and B is any sequence of nodes all connected by arcs to get you from A to B. The path is blocked if (now with arrows returned) there is a collider on the path. [A variable C is a collider on a path if the arcs on the path that meet at C both have arrows pointing at C.] If a path is not blocked, we say it is unblocked, active, or open, all of which are synonymous.

Then $(A \amalg B)_G$ if and only if every path from A to B is blocked. If even one path is unblocked, we write $(A \not\amalg B)_G$.

We say two variables (A_1, A_2) are d-separated from two other variables (B_1, B_2), i.e., $((A_1, A_2) \amalg (B_1, B_2))_G$ if and only if all paths between any variable A_j and any variable B_ℓ are blocked. That is, you simply test d-separation between each variable in the second group and each variable in the first group. If there are two variables in the first group and two in the last group, you have to do four separate checks of d-separation.

Conditional independence: We say two variables A and B are d-separated given (or by) a set of variables $Z = (Z_1, \ldots, Z_k)$ if all paths between A and B are blocked where, when we can condition on Z, a path between A and B is blocked if *(i)* there is any variable $Z_m \in Z$ on the path that is not a collider or *(ii)* there is a collider on the path such that neither the collider itself nor any of its descendants are in Z. Generalizing this idea to d-separation of (A_1, A_2) from (B_1, B_2) given Z is just as above. That is, we check that each variable in the first set is d-separated from each variable in the second set conditional on Z.

Statistical interpretation of d-separation: As we have seen, a DAG represents a statistical model. That is, it represents a set of distributions whose density can be factorized as the product of the probability of each variable given its parents. Suppose we wish to know whether, for all distributions in the model (i.e., all distributions represented by our DAG), whether $A \amalg B \mid Z$. The answer to this question is $A \amalg B \mid Z$ for all distributions represented by the model if and only if A is d-separated from B by variables Z, i.e., $(A \amalg B \mid Z)_G$.

Suppose now that A is not d-separated from B given Z on the graph. Then there exists at least one distribution, represented by the DAG, for which A and B are dependent given Z. [Note that there may be other distributions in the model (represented by the DAG) for which A and B are independent given Z.] For example, if Z is the empty set and our DAG is simply $A \rightarrow B$, then clearly A is not d-separated from B; yet distributions in which A is independent of B are represented by our DAG since our DAG is complete (all arrows are present), and thus the DAG represents all distributions for (A, B).

The G-computation algorithm formula: Given a DAG G with variables $(A, B, C, D, \ldots, Z)$, suppose the DAG is complete and the ordering of the DAG is alphabetical, that is, the arrow between two variables has its head pointing at the variable later in the alphabet. Suppose we want to figure out the density of all the uncontrolled variables in a hypothetical study where we intervene and set variables C, O, and X to c,o, x. The joint density of all the other variables in this hypothetical study is derived as follows. Write the joint density of all the observed variables on the graph using the usual DAG factorization, i.e., the product over the 26 variables of $f\,(\mathit{variable} \mid \mathit{parent})$. [Note Z has 25 parents while A has no parents.] Next remove from this product the terms corresponding to the densities of the set (manipulated) variables given their parents. That is, $f\,(c \mid \mathit{parents}\ C)\,, f\,(o \mid \mathit{parents}\ O)\,, f\,(x \mid \mathit{parents}\ X)$. Finally, be sure that in the remaining terms, the value you have for the variables O, C, and X are the values you set these variables to. Finally, to get the marginal distribution of any single variable, say, Y in the study where you have intervened and set C, O, and X to particular values C, O, X of interest, you take the previous joint density and sum out (integrate out) all the remaining 22 variables [i.e., all variables but Y and the variables O, C, and X that you have set]. Note the G-computation algorithm formula is defined in terms of the distribution of the observed variables that you actually see in your study and can in principle be consistently estimated (by counting) in a sufficiently large study. The formula has the distribution you would see if you intervene and set C, O, and X to particular values of interest if the original DAG is a causal DAG — i.e., there are no other unmeasured confounders in the world that you need to worry about. Of course, in an observational study, you don't know whether this is true. In a randomized study (where there was physical randomization), you should know whether this was true because you know which variables were used to assign the treatments C, O, and X.

(a) (b)

Figure 1: (a) A causal DAG representing two unconfounded variables; (b) A causal DAG representing the presence of confounding.

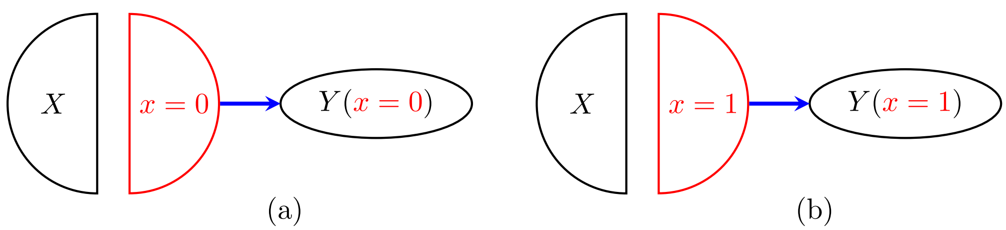

Figure 2: The graphs resulting from splitting node X in the graph in Figure 1(a), and intervening to set a particular value. (a) setting X to 0; (b) setting X to 1.

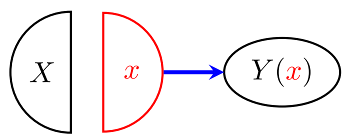

Figure 3: A template representing the two graphs in Figure 2.

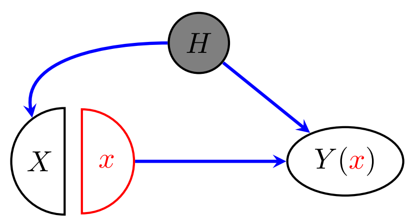

Figure 4: The template resulting from intervening on X in the graph in Figure 1(b).

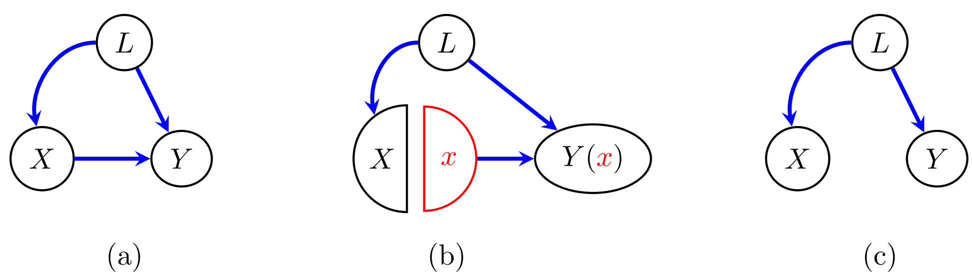

Figure 5: Adjusting for confounding. (a) The original causal graph. (b) The template $\mathcal{G}(x)$, which shows that $Y(x) \perp\!\!\!\perp X \mid L$. (c) The DAG $\mathcal{G}_{\underline{X}}$ obtained by removing edges from X advocated in Pearl (1995, 2000, 2009) to check his 'backdoor condition'.

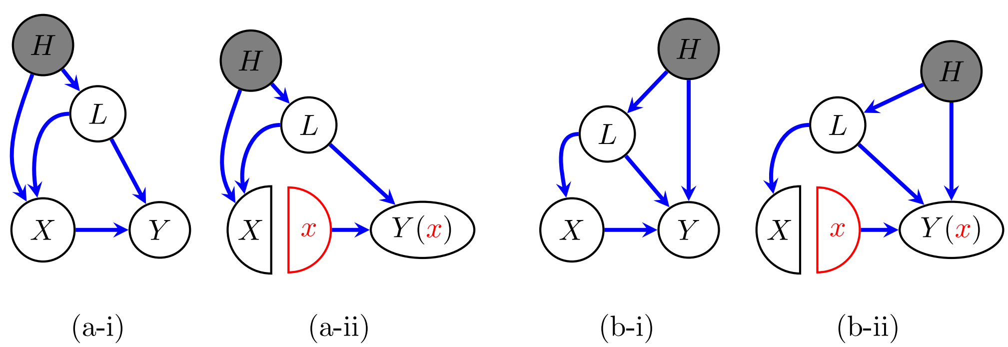

Figure 6: Further examples of adjusting for confounding. (a-i) A graph $\mathcal{G}$; (a-ii) the template $\mathcal{G}(x)$; (b-i) A graph $\mathcal{G}'$; (b-ii) the template $\mathcal{G}'(x)$. H is an unobserved variable in $\mathcal{G}$ and $\mathcal{G}'$. Both SCoTs imply $Y(x) \perp\!\!\!\perp X \mid L$.

(1) *Split Nodes*: For every $A \in \mathbf{A}$ split the node into a random and fixed component, labelled A and a respectively, as follows:

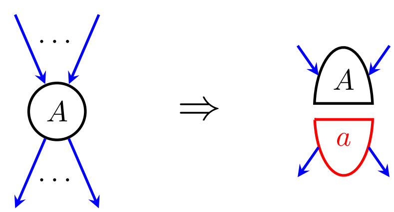

Splitting: Schematic Illustrating the Splitting of Node A

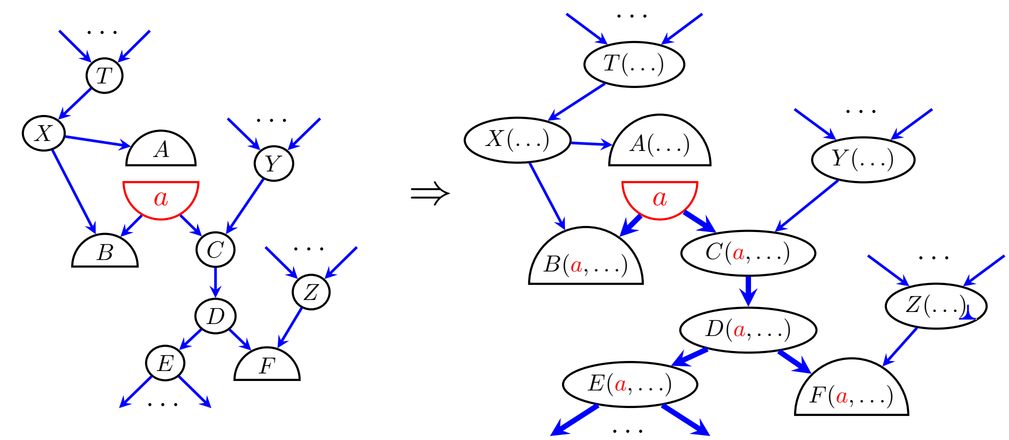

Labelling: Schematic showing the nodes $V(\mathbf{a}_V)$ in $\mathcal{G}(\mathbf{a})$ for which $a \in \mathbf{a}_V^{\cap}$

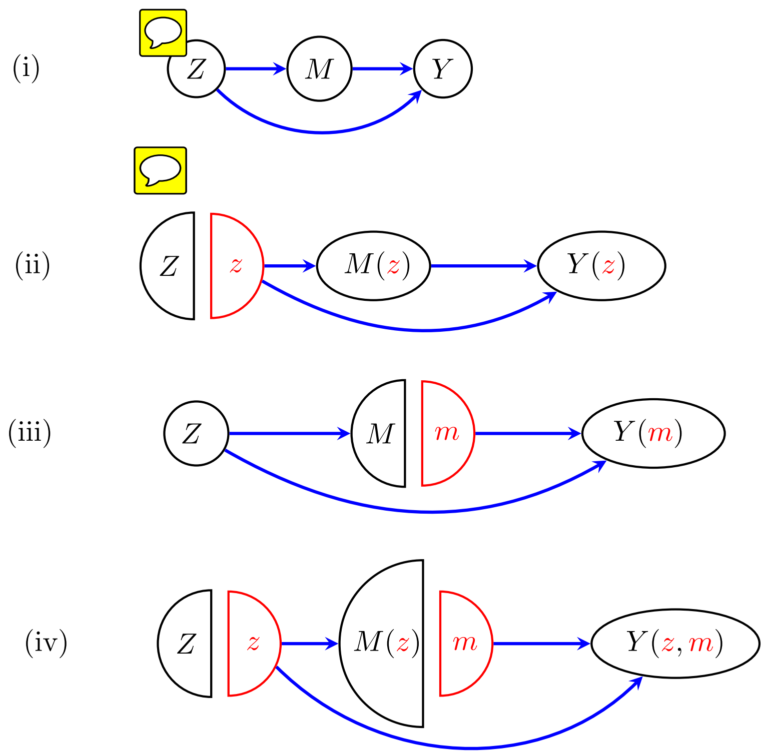

Figure 7: (i) A DAG $\mathcal{G}$ with treatment (Z), mediator (M) and response (Y) in the absence of confounding. minimal templates: (ii) $\mathcal{G}(z)$; (iii) $\mathcal{G}(m)$; (iv) $\mathcal{G}(m, z)$.

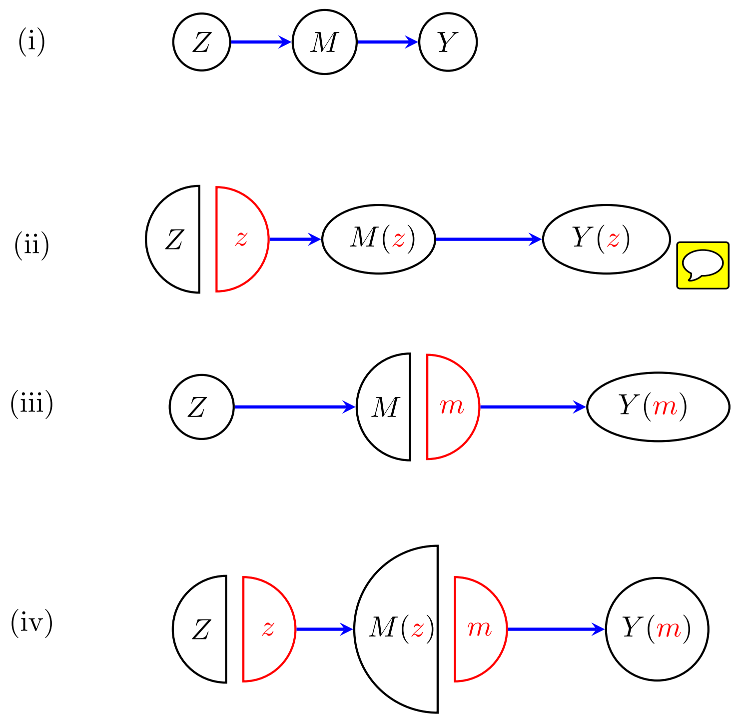

Figure 8: (i) The DAG $\mathcal{G}$ from Figure 7 under the additional assumption that there is no direct effect of treatment (Z) on the response (Y). Templates: (ii) $\mathcal{G}(z)$; (iii) $\mathcal{G}(m)$; (iv) $\mathcal{G}(z, m)$.

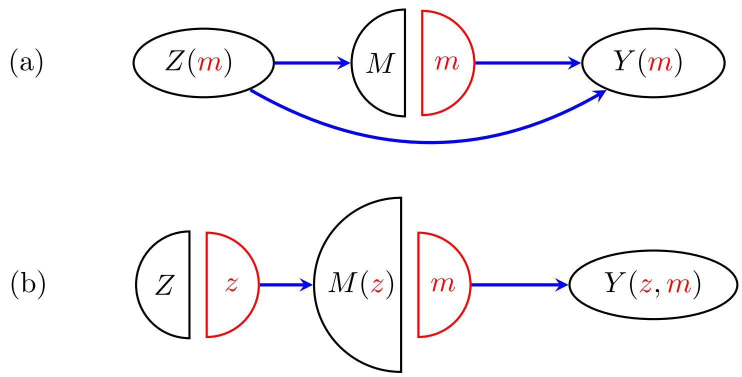

Figure 9: Non-minimal SCoTs: (a) A template $\mathcal{G}'(m)$ formed from the minimal SCoT $\mathcal{G}(m)$ in Figure 7(iii) by adding m to Z; A template (b) $\mathcal{G}'(z, m)$ is formed from $\mathcal{G}(z, m)$ in Figure 8(iv) by adding z to $Y(m)$.

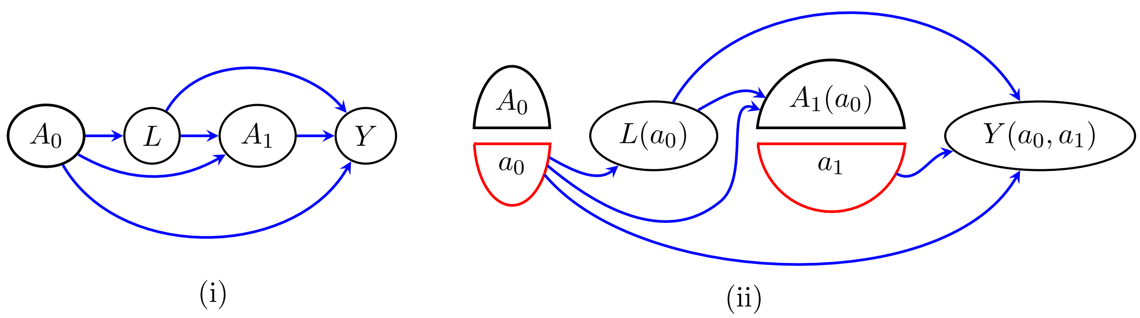

Figure 10: (i) (i) A complete DAG $\mathcal{G}$; (ii) the template $\mathcal{G}(a_0, a_1)$.

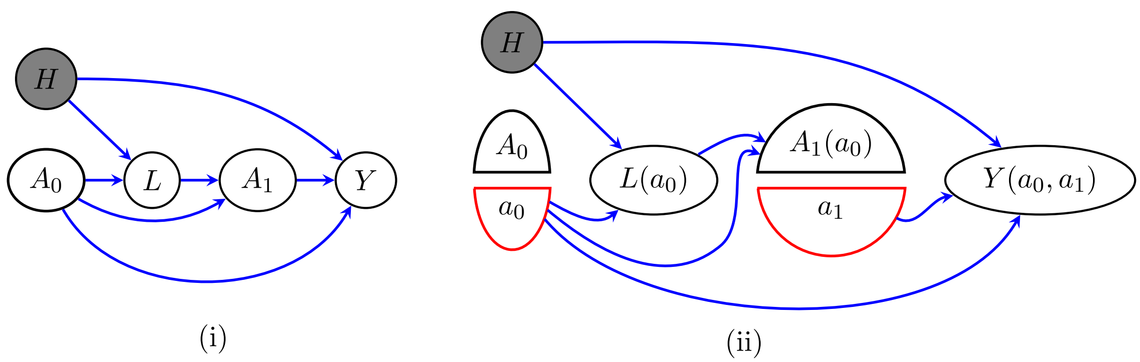

Figure 11: (i) (i) A DAG $\mathcal{G}$ describing a sequentially randomized trial; (ii) the template $\mathcal{G}(a_0, a_1)$.

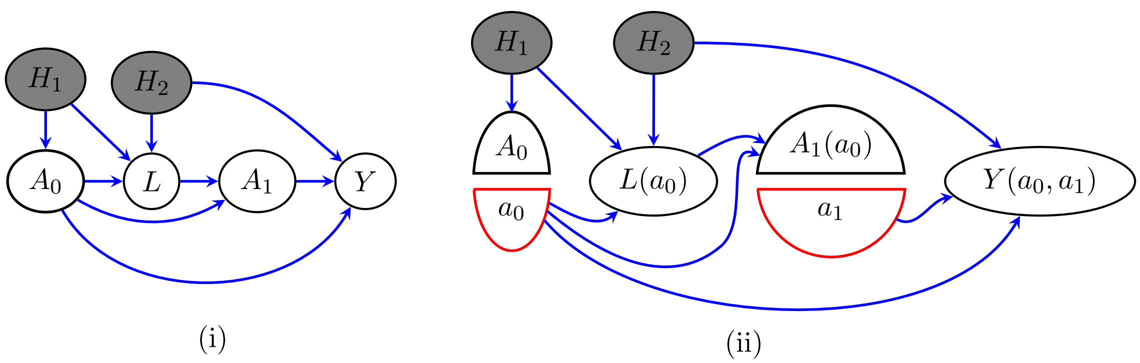

Figure 12: (i) A DAG $\mathcal{G}$ in which initial treatment is confounded, while the second treatment is sequentially randomized; (ii) the SCoT $\mathcal{G}(a_0, a_1)$. L is known to have no direct effect on Y, except indirectly via the effect on A_1

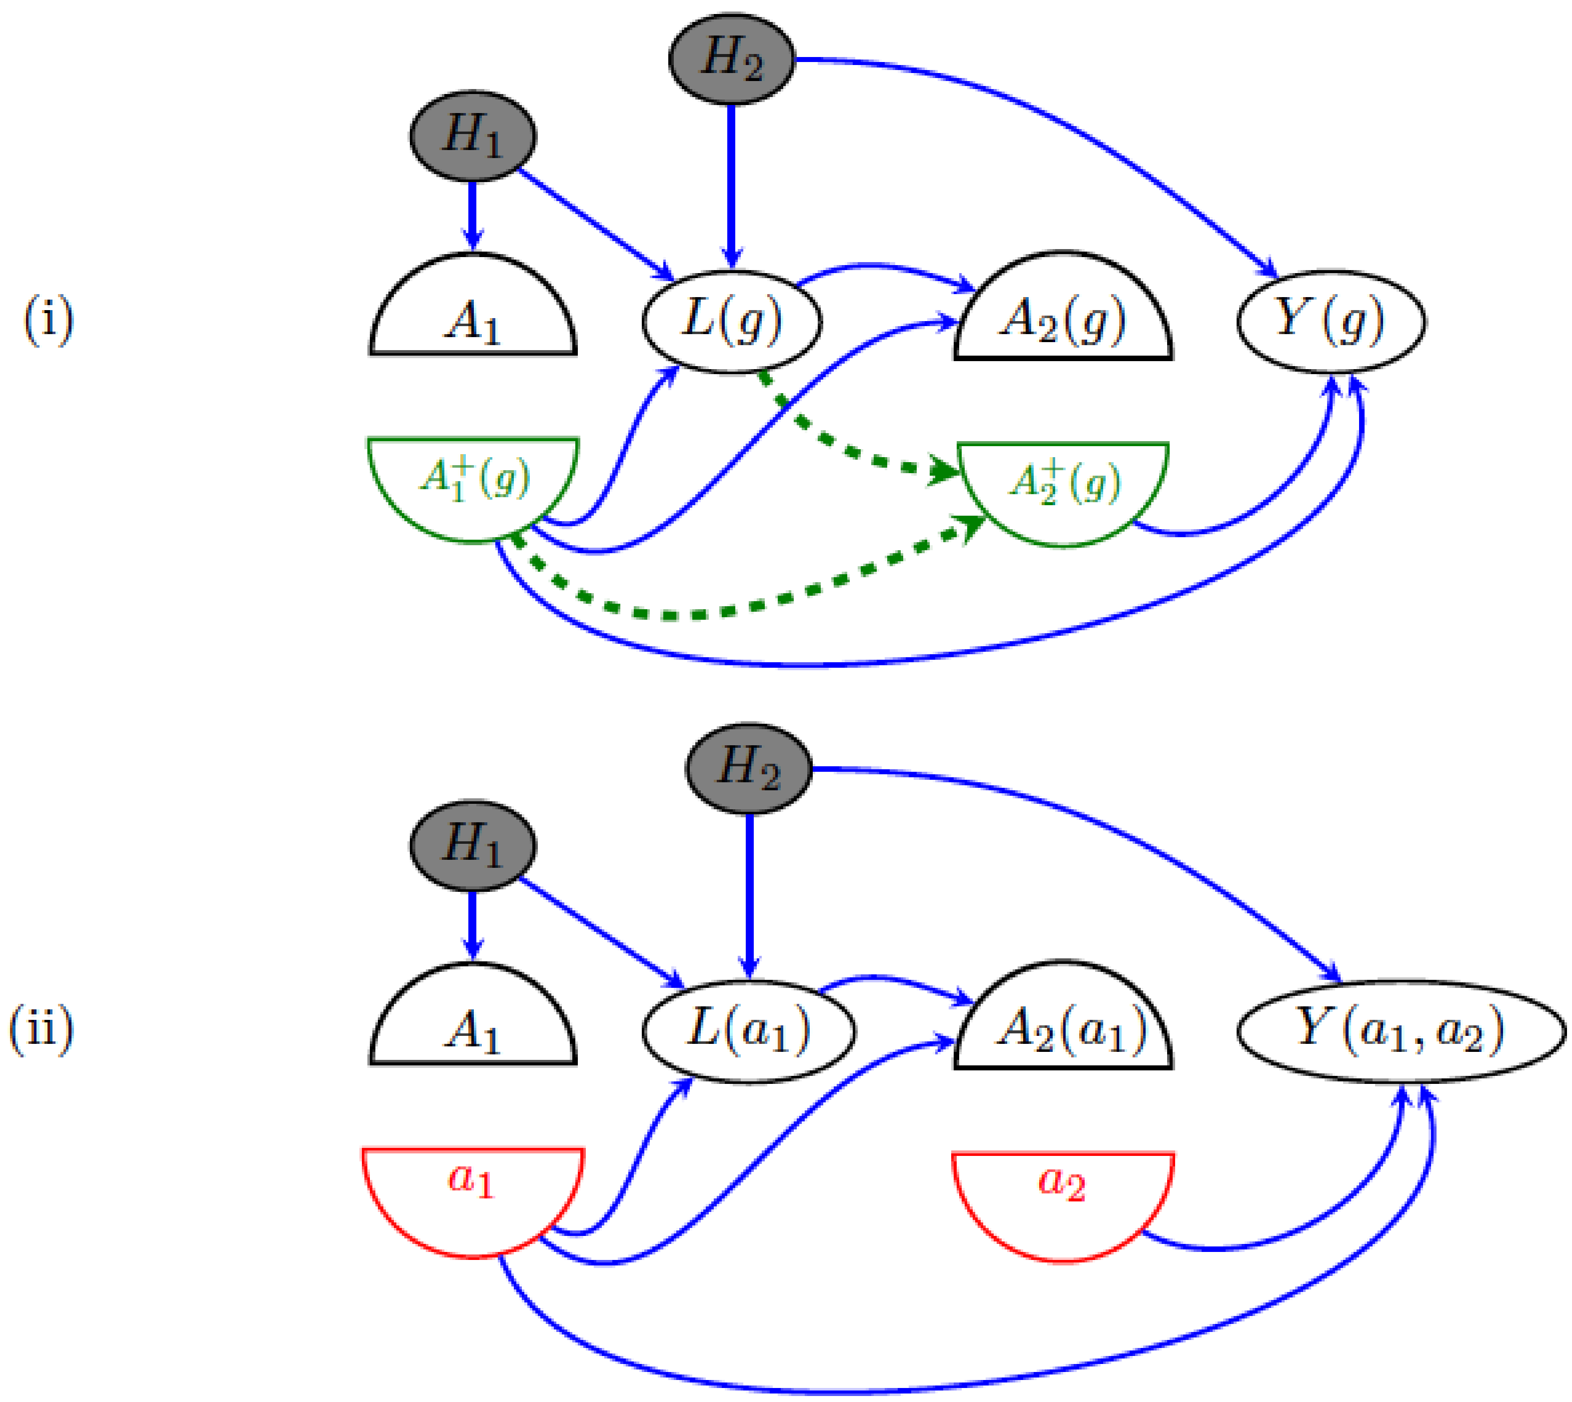

Figure 22: (i) the dynamic SWIG $\mathcal{G}(g)$ derived from the template in Figure 13(ii) under a dynamic regime g for which treatment at the second time $A_2^+(g)$ depends on past treatment and covariate history ($A_1^+(g)$ and $L(g)$). $P(Y(g))$ is not identified since $A_2^+(g)$ is a function of $L(g)$ and hence $L(g) \in \mathbb{Q}_0(g)$, but there is a d-connecting path from A_1 to $L(a_1)$ in $\mathcal{G}(a_1, a_2)$, as shown in (ii). Thus the d-separation relation (50) in Theorem 27 does not hold. Since g here does not depend on A_1 or $A_2(g)$ we may also apply Theorem 28 to $\mathcal{G}(g)$ directly: Since there is a path d-connecting A_1 and $Y(g)$ (given $\emptyset$) condition (iii) fails to hold for $k = 1$.

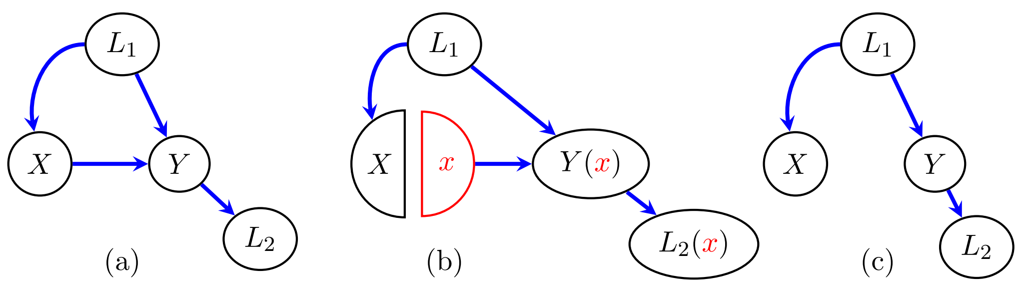

Figure 13: Simplification of the backdoor criterion. (a) The original causal graph $\mathcal{G}$. (b) The template $\mathcal{G}(x)$, which shows that $Y(x) \perp\!\!\!\perp X \mid L_1$, but does not imply $Y(x) \perp\!\!\!\perp X \mid \{L_1, L_2\}$. (c) The DAG $\mathcal{G}_{\underline{X}}$ obtained by removing edges from X advocated in Pearl (2000, 2009).

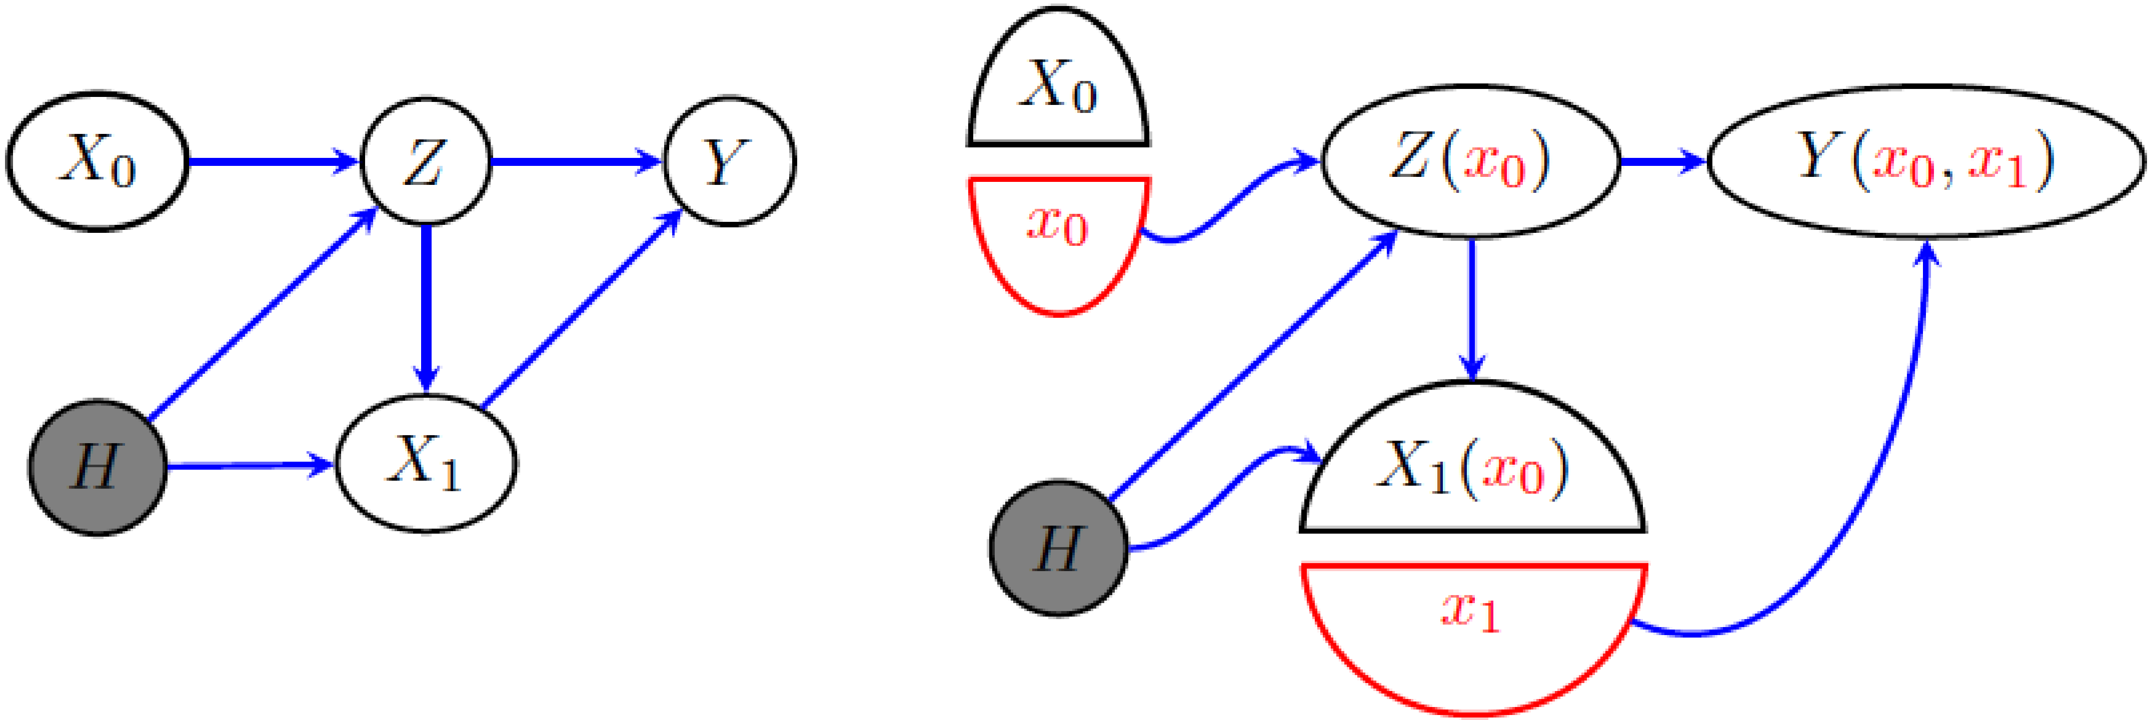

(a) (b)

Figure 15: (a) A DAG $\mathcal{G}$ in which Robins' condition (38) is claimed not to hold; see Ex. 11.3.3, Fig. 11.12 in Pearl (2009, p.353); H is unobserved ; (b) the template $\mathcal{G}(x_0, x_1)$.

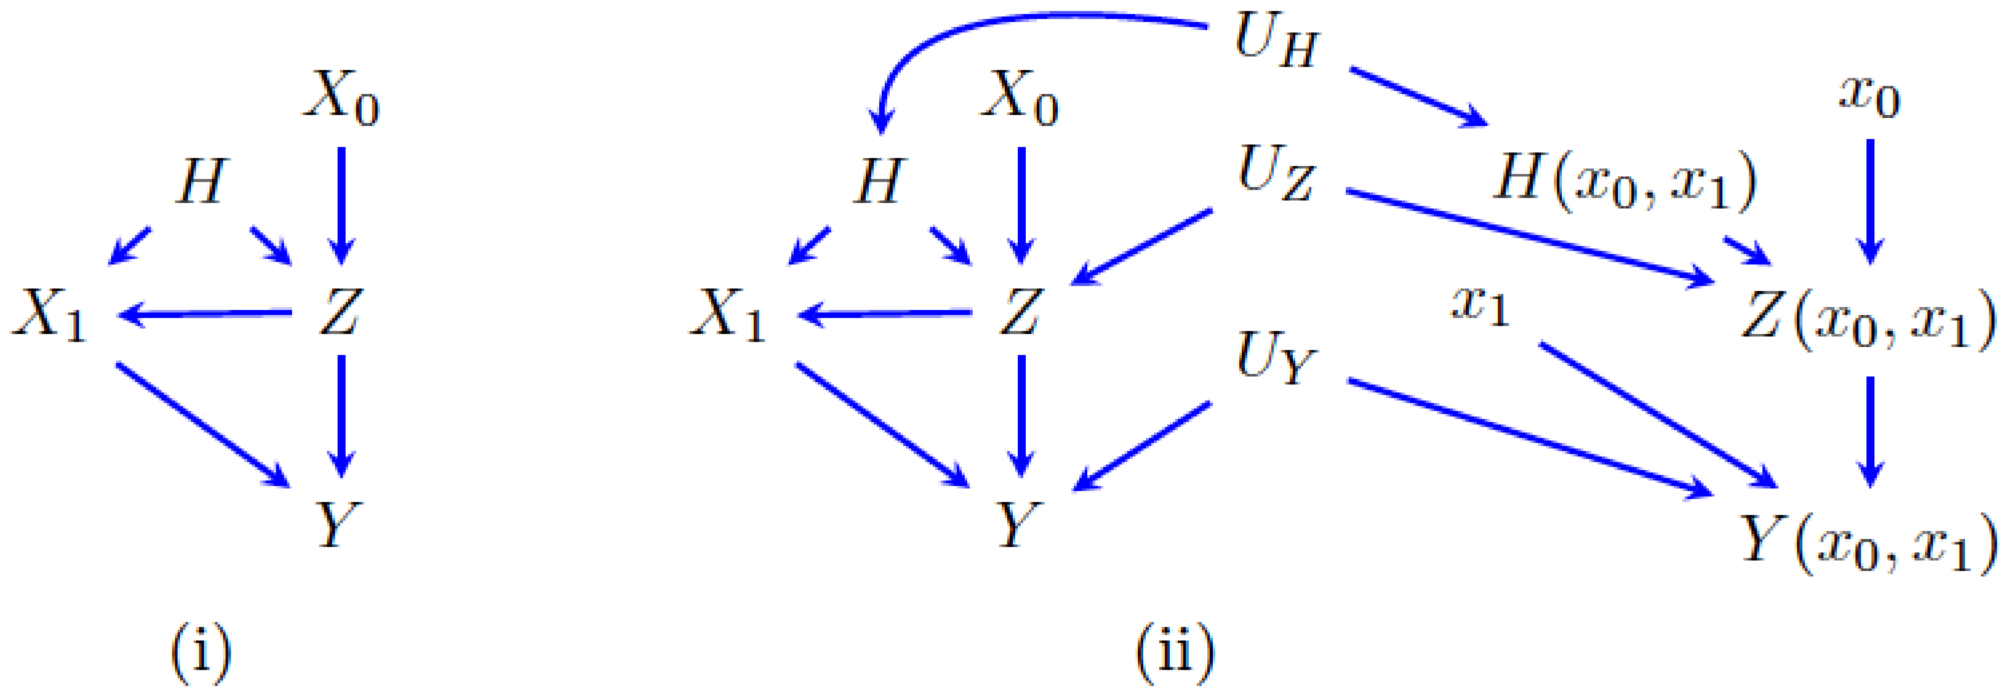

Figure 28: Failure of the simple twin network method of Pearl (2000, 2009) in Pearl's Example 11.3.3. (i) The original DAG; (ii) The twin network after intervention on X_0 and X_1. The twin network fails to reveal that $Y(x_0, x_1) \perp\!\!\!\perp X_1 \mid Z, X_0 = x_0$. This 'extra' independence holds in spite of d-connection because (by consistency) when $X_0 = x_0$, then $Z = Z(x_0) = Z(x_0, x_1)$. (Note that $Y(x_0, x_1) \not\!\perp\!\!\!\perp X_1 \mid Z, X_0 \neq x_0$.)