11.3.4 Data vs. Knowledge in Covariate Selection
(From Causality, Second edition, 2009.)

Scroll down to beginning of section.

11.3.4 Data vs. Knowledge in Covariate Selection

What then can be done in the absence of a causal graph? One way is to postulate a plausible graph, based on one's understanding of the domain, and check if the data refutes any of the statistical claims implied by that graph. In our case, the graph of Figure 11.8(b) advertises several such claims, cast as conditional independence constraints, each associated with a missing arrow in the graph:

$$\begin{array}{lll} V \perp\!\!\!\perp X \mid Z_1, W_1 & V \perp\!\!\!\perp W_2 & Z_2 \perp\!\!\!\perp W_1 \mid W_2 \\ V \perp\!\!\!\perp Y \mid X, Z_2, W_2 & Z_1 \perp\!\!\!\perp Z_2 \mid V, W_1, W_2 & Z_1 \perp\!\!\!\perp W_2 \mid W_1 \\ V \perp\!\!\!\perp W_1 & X \perp\!\!\!\perp Z_2 \mid Z, W_1, W_2 & X \perp\!\!\!\perp \{V, Z_2\} \mid Z_1, W_1, W_2. \end{array}$$

Satisfying these constraints does not establish, of course, the validity of the causal model postulated because, as we have seen in Chapter 2, alternative models may exist which satisfy the same independence constraints yet embody markedly different causal structures, hence, markedly different admissible sets and effect estimands. A trivial example would be a complete graph, with arbitrary orientation of arrows which, with a clever choice of parameters, can emulate any other graph. A less trivial example, one that is not sensitive to choice of parameters, lies in the class of equivalent structures, in

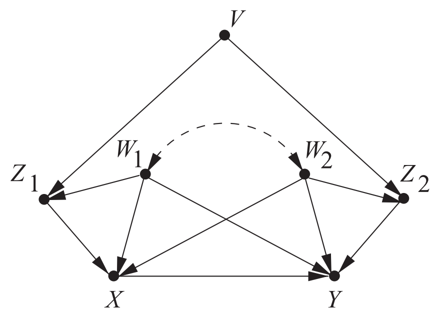

Figure 11.9 A model that is dependence-wise indistinguishable from that of Figure 11.8 (b), in which the irreducible sets $\{Z_1, W_1, W_2\}$, $\{W_1, W_2, V\}$, and $\{W_1, W_2, Z_2\}$ are admissible.

which all conditional independencies emanate from graph separations. The search techniques developed in Chapter 2 provide systematic ways of representing all equivalent models compatible with a given set of conditional independence relations.

For example, the model depicted in Figure 11.9 is indistinguishable from that of Figure 11.8(b), in that it satisfies all the conditional independencies implied by the latter, and no others.[6] However, in contrast to Figure 11.8(b), the sets $\{Z_1, W_1, W_2\}$, $\{V, W_1, W_2\}$, and $\{Z_2, W_1, W_2\}$ are admissible. Adjusting for the latter would remove bias if the correct model is Figure 11.9 and might produce bias if the correct model is Figure 11.8(b).

Is there a way of telling the two models apart? Although the notion of "observational equivalence" precludes discrimination by statistical means, substantive causal knowledge may provide discriminating information. For example, the model of Figure 11.9 can be ruled out if we have good reasons to believe that variable W_2 cannot have any influence on X (e.g., it may occur *later* than X,) or that W_1 could not possibly have direct effect on Y.

The power of graphs lies in offering investigators a transparent language to reason about, to discuss the plausibility of such assumptions and, when consensus is not reached, to isolate differences of opinion and identify what additional observations would be needed to resolve differences. This facility is lacking in the potential-outcome approach where, for most investigators, "strong ignorability" remains a mystical black box.

In addition to serving as carriers of substantive judgments, graphs also offer one the ability to reject large classes of models without testing each member of the class. For example, all models in which V and W_1 are the sole parents of X, thus rendering $\{V, W_1\}$ (as well as C) admissible, could be rejected at once if the condition $X \perp\!\!\!\perp Z_1 \mid V, W_1$ does not hold in the data.

In Chapter 3, for example, we demonstrated how the measurement of an additional variable, mediating between X and Y, was sufficient for identifying the causal effect of X on Y. This facility can also be demonstrated in Figure 11.8(b); measurement of a variable Z judged to be on the pathway between X and Y would render $P(y \mid do(x))$ identifiable and estimable through equation (3.29). This is predicated, of course, on Figure 11.8(b) being the correct data-generating model. If, on the other hand, it is Figure 11.9 that represents the correct model, the causal effect would be given by

$$\begin{aligned} P(y \mid do(x)) &= \textstyle\sum_{pa_X} P(y \mid pa_X, x) P(pa_X) \\ &= \textstyle\sum_{z_1, w_1, w_2} P(y \mid x, z_1, w_1, w_2)\, P(z_1, w_1, w_2), \end{aligned}$$

[6] Semi-Markovian models may also be distinguished by functional relationships that are not expressible as conditional independencies (Verma and Pearl 1990; Tian and Pearl 2002b; Shpitser and Pearl 2008). We do not consider these useful constraints in this example.

which might or might not agree with equation (3.29). In the latter case, we would have good reason to reject the model in Figure 11.9 as inconsistent, and seek perhaps additional measurements to confirm or refute Figure 11.8(b).

Auxiliary experiments may offer an even more powerful discriminatory tool than auxiliary observations. Consider variable W_1 in Figure 11.8(b). If we could conduct a controlled experiment with W_1 randomized, instead of X, the data obtained would enable us to estimate the causal effect of X on Y with no bias (see Section 3.4.4). At the very least, we would be able to discern whether W_1 is a parent of X, as in Figure 11.9, or an indirect ancestor of X, as in Figure 11.8(b).

In an attempt to adhere to traditional statistical methodology, some causal analysts have adopted a method called "sensitivity analysis" (e.g., Rosenbaum 2002, pp. 105–170), which gives the impression that causal assumptions are not invoked in the analysis. This, of course, is an illusion. Instead of drawing inferences by assuming the absence of certain causal relationships in the model, the analyst tries such assumptions and evaluates how strong alternative causal relationships must be in order to explain the observed data. The result is then submitted to a judgment of plausibility, the nature of which is no different from the judgments invoked in positing a model like the one in Figure 11.9. In its richer setting, sensitivity analysis amounts to loading a diagram with causal relationships whose strength is limited by plausibility judgments and, given the data, attempting to draw conclusions without violating those plausibility constraints. It is a noble endeavor, which thus far has been limited to problems with a very small number of variables. The advent of diagrams promises to expand the applicability of this method to more realistic problems.