TECHNICAL REPORT
R-449
November 2015

Exogeneity and Robustness

Bryant Chen and Judea Pearl
University of California, Los Angeles
Computer Science Department
Los Angeles, CA, 90095-1596, USA
(310) 825-3243

November 4, 2015

Abstract

A common practice for detecting misspecification is to perform a "robustness test", where the researcher examines how a regression coefficient of interest behaves when variables are added to the regression. Robustness of the regression coefficient is taken as evidence of structural validity. However, there are numerous pitfalls that can befall a researcher when performing such tests. For example, we demonstrate that certain regressors, when added to the regression, will induce a shift in the coefficient of interest even when structurally valid. Such robustness tests would produce false alarm, suggesting that the model is misspecified when it is not. For a robustness test to be informative, the variables added to the regression must be carefully chosen based on the model structure. We provide a simple criterion that allows researchers to quickly determine which variables, when added to the regression, constitute informative robustness tests. We also explore the extent to which robustness tests are able to detect bias, demonstrating that robustness tests enable detection of bias due not only to omitted observable variables but omitted unobservable variables as well. Finally, we give two extended examples using simulated data to demonstrate how the material in this paper can be used to conduct informative robustness tests.

1 Introduction

Suppose that economic theory dictates the structural equation

$$Y = \beta_0 D + \alpha_0 Z + U, \tag{1}$$

where D is an observable cause of interest, Y is an outcome of interest, Z is an observable driver of Y, and U represents unobservable drivers of Y, often called a "disturbance" or "error" term[1]. To estimate the value of β_0, we may regress Y on D, and

[1] D and Z may also be vectors of variables without changing any of the results in this paper but for convenience we assume that they are singletons.

a set of variables, $W = \{W_1, W_2, ..., W_k\}$, where $Z \subseteq W$. This gives the regression equation

$$Y = r_0 D + r_1 W_1 + r_2 W_2 + ... + r_k W_k + \epsilon, \qquad (2)$$

where ϵ is the regression residual and r_0 is the partial regression coefficient of D when Y is regressed on D and W, that is, $r_0 = r_{YD.W}$.

It is well known that β_0 is identified and equal to $r_{YD.W}$ if W satisfies the conditional independence $D \perp\!\!\!\perp U | W$, often called the *conditional exogeneity* assumption. If this assumption fails then the regression coefficient $r_{YD.W}$ will generally not equal β_0. Unfortunately, since U is unobservable, it is impossible to test whether this assumption holds, and the task of defending the validity of $D \perp\!\!\!\perp U | W$ is delegated to human judgment, which is vulnerable to two sources of error. The first, of course, is model misspecification: although $D \perp\!\!\!\perp U | W$ can in principle be verified from the hypothesized structure, one can never be sure whether these theoretical assumptions are valid. Second, to judge whether $D \perp\!\!\!\perp U | W$ holds in a given specified model can be formidable when multiple equations and multiple U factors are present, some correlated with the observable variables and some with other U factors.

While the conditional exogeneity assumption is not directly testable, one can identify various implications of the model structure that can be used to test it against data. If these implications are found compatible with the data then the model gains credibility. Consequently, if the model implies conditional exogeneity then it also gains credibility. Testing the conditional exogeneity assumption along these lines requires two steps. First, to identify and test all testable implications of the model, and second, to verify that conditional exogeneity holds in the model. The first step can be rather involved and is rarely performed in practice. Instead, practitioners resort to shortcuts– testing only a subset of implications deemed relevant to the conditional exogeneity assumption[2].

A common exercise in empirical studies, which utilizes such shortcuts, is to check the "robustness" of certain regression coefficients when the regression specification is modified by including or excluding "control" covariates. Movement by the regression coefficients of D is then taken as evidence of omitted variable bias or misspecification. In a recent survey of non-experimental empirical work, Lu and White (2014) found that of the 76 papers involving data analysis published in *The American Economic Review* during 2009, 23 perform a robustness check along the lines just described. Similarly, Oster (2013) found that 75% of 2012 papers published in *The American Economic Review*, *Journal of Political Economy*, and *Quarterly Journal of Economics* explored the sensitivity of results to varying control sets[3].

The intuition behind this procedure is rooted on the following heuristic: if bias is caused by some set W of confounders, then controlling for W (by adding W to the regression equation) should eliminate that bias. Any further control would then be unnecessary, and should leave r_0 unaltered. The invariance of r_0 to additional regressors is taken as evidence that all confounders have been accounted for, and none

[2]The methods for testing conditional exogeneity proposed by White and Chalak (2010) can be understood in this manner. However, these methods require specific structural assumptions beyond conditional exogeneity and are, therefore, not applicable to all models.

[3]The practice of assessing missing variable bias by observing the sensitivity of an estimator to additional controls is not unique to economics. A recent survey of articles in major epidemiologic journals by Walter and Tiemeier (2009) found that 15% of papers used a "change-in-estimate" criterion to select covariates.

remain outside the set W. Conversely, if r_0 shifts with the addition of regressors beyond W, it is taken as evidence that W was not sufficient to cover all confounders and, consequently, r_0 is likely to be biased and the conditional exogeneity $D \perp\!\!\!\perp U | W$ is likely violated.

This heuristic assumes that bias reduction is monotonic with the number of regressors, that is, that adding regressors cannot create bias where none exists. We will show that this assumption is false. Certain regressors, if added to W will necessarily change r_0, even when W is sufficient to satisfy the exogeneity condition $D \perp\!\!\!\perp U | W$. We call such regressors *shift-producing*. Robustness tests involving shift-producing covariates are non-informative and produce false alarm when exogeneity holds. We will show how shift-producing regressors can be identified from the model's structure.

For informative tests, the connection between robustness and exogeneity is as follows. If conditional exogeneity holds before and after the addition of regressors then obviously r_0 will be invariant. Conversely, if r_0 shifts then conditional exogeneity must be violated before or after the addition. As a result, if the hypothesized model structure dictates that exogeneity holds in both cases then the added regressor may be informative since failure of the corresponding robustness test indicates model misspecification. However, as mentioned previously, the task of determining whether exogeneity holds, even in a well-specified model, can be formidable. As formal and transparent representations of the model structure, causal graphs provide researchers with the means to determine, by inspection, whether a given set of variables satisfies the conditional exogeneity assumption. Utilizing a graphical condition, called *single-door criterion*, we are able to quickly identify sets, W_1 and W_2, for which the model implies $r_{YD.W_1} = r_{YD.W_1W_2} = \beta_0$. We will show that comparing the corresponding regressions, $r_{YD.W_1}$ and $r_{YD.W_1W_2}$, constitutes an informative robustness test in that failure both implies model misspecification and is possible when the conditional exogeneity assumption is violated. As a result, researchers can focus on the economic plausibility of the model structure and substitute all other judgments with sound and reliable mechanical procedures when identifying β_0 and finding informative covariates for robustness testing.

Graphs have been utilized by economists to communicate causal structure and facilitate economic problems since the 1930s (Tinbergen, 1939). Orcutt (1952), for instance, used graphs to represent possible causal structures for a given set of variables and gave examples illustrating that some graphs are incompatible with certain conditional independences among the variables (e.g. zero partial correlation between Z and X given Y is incompatible with the chain $Y \rightarrow X \rightarrow Z$), thus allowing researchers to reject certain causal structures. Tinbergen (1939), Wold (1954), and other practitioners of process analysis also employed graphs to convey the causal relationship between variables (Hoover, 2004). More recently, graphs have been utilized by White and Chalak (2009), White and Lu (2011), and Hoover and Phiromswad (2013) to facilitate problems of identification, optimization, identifying instrumental variables from data, and more. We will use graphical models primarily to detect conditional independences and verify identifying assumptions.

Our paper is structured in the following way: We will begin Section 2 by introducing graphical representations of structural models with special interest on graph separation, a notion that will play a pivotal role in the results that follow. In Section 3, we introduce the single-door criterion, a necessary and sufficient condition for the

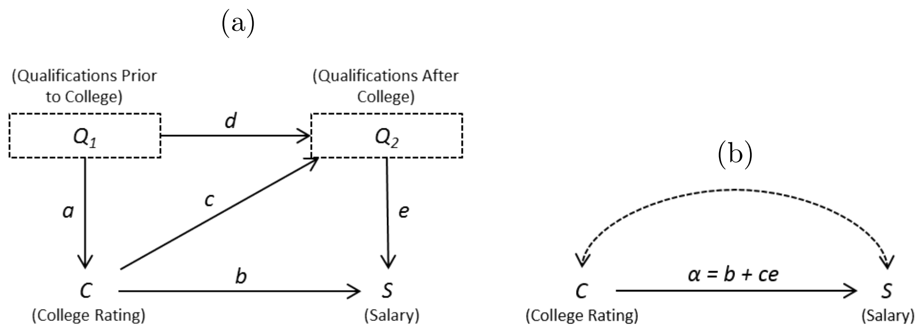

Figure 1: (a) Model with latent variables (Q_1 and Q_2) shown explicitly (b) Same model with latent variables summarized

identification of β_0 using regression that will allow us to characterize shift-producing covariates and informative robustness tests in Section 4. We also explore the extent to which robustness tests are able to detect bias in Section 4, demonstrating that robustness tests are to detect not only omitted observable variables but omitted unobservable variables as well. In Section 5, we will adapt the results in Sections 3 and 4 for situations where we are interested in the total effect of D on Y, rather than the direct effect. Finally, in Section 6, we will give two extended examples using simulated data to demonstrate how the material in this paper can be used to avoid shift-producing regressors and conduct informative robustness tests.

2 Preliminaries

2.1 Causal Graphs

We introduce causal graphs by way of example. Suppose we wish to estimate the effect of attending an elite college on future earnings. Clearly, simply regressing earnings on college rating will not give an unbiased estimate of the target effect. This is because elite colleges are highly selective so students attending them are likely to have qualifications for high-earning jobs prior to attending the school. This background knowledge can be expressed in the following model specification.

Model 1.

$$
\begin{aligned}
Q_1 &= U_1 \\
C &= a \cdot Q_1 + U_2 \\
Q_2 &= c \cdot C + d \cdot Q_1 + U_3 \\
S &= b \cdot C + e \cdot Q_2 + U_4,
\end{aligned}
$$

where Q_1 represents the individual's qualifications prior to college, Q_2 represents qualifications after college, C contains attributes representing the quality of the college attended, and S the individual's salary. When specifying models throughout the paper, the error terms (U variables) are independent unless otherwise stated. In this case, the model assumes that U_1, U_2, U_3, and U_4 are independent of one another.

(a)

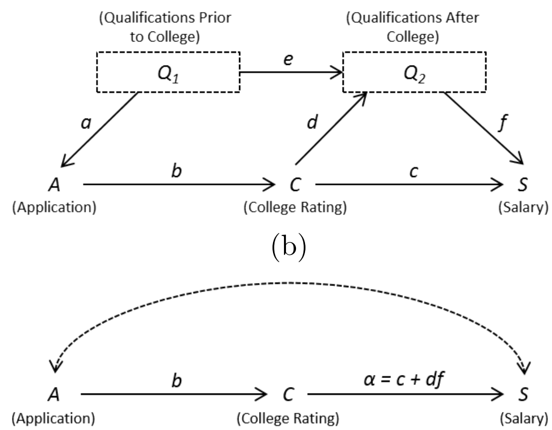

Figure 2: Graphs associated with Model 3 in the text (a) with latent variables shown explicitly (b) with latent variables summarized

Figure 1a is a causal graph that represents this model specification. Each variable in the model has a corresponding node or vertex in the graph. Additionally, for each equation, arrows are drawn from the independent variables to the dependent variables. These arrows reflect the direction of causation. In some cases, we may label the arrow with its corresponding structural coefficient as in Figure 1a. Error terms are typically not displayed in the graph.

If Q_1 and Q_2 are unobservable or latent variables their influence on S is generally attributed to S's error term. By removing them, we obtain the following model specification, where $\sigma_{U_C U_S}$ is the covariance of U_C and U_S:

Model 2.

$$\begin{aligned} C &= U_C \\ S &= \beta C + U_S \\ \sigma_{U_C U_S} &\neq 0 \end{aligned}$$

In this case, $U_C = a \cdot Q_1 + U_1$ and $U_S = eQ_2 + U_4$, and the background information specified by Model 1 imply that the error term of S, U_S, is correlated with U_C and, therefore, correlated with C. As a result, exogeneity does not hold. Dependence between error terms is depicted in the causal graph as a bidirected arc between the variables whose error terms are dependent as in Figure 1b.

It is clear that β is not identified in Model 2. However, if we include the strength of an individual's college application, A, as shown in Figure 2a, we obtain the following model:

Model 3.

$$\begin{aligned}
Q_1 &= U_1 \\
A &= a \cdot Q_1 + U_2 \\
C &= b \cdot A + U_3 \\
Q_2 &= e \cdot Q_1 + d \cdot C + U_4 \\
S &= c \cdot C + f \cdot Q_2 + U_5.
\end{aligned}$$

By removing the latent variables from the model specification we obtain:

Model 4.

$$\begin{aligned}
A &= a \cdot Q_1 + U_A \\
C &= b \cdot A + U_C \\
S &= \beta \cdot C + U_S,
\end{aligned}$$

Now, β is identified and can be estimated using the regression of S on C and A. While this conclusion can, in principle, also be reached by consulting the equations themselves, in more complex models doing so can be infeasible. In Section 3, we will demonstrate how to obtain this and more complicated identification results using the causal graph.

In summary, the causal graph is constructed from the model equations in the following way: Each variable in the model has a corresponding vertex or node in the graph. For each equation, arrows are drawn in the graph from the dependent variables to the independent variable. Finally, if the error terms of any two variables are dependent, then a bidirected edge is drawn between the two variables.

Before continuing, we review some basic graph terminology. If an arrow, called (X, Y), exists from X to Y we say that X is a *parent* of Y. If there exists a sequence of arrows all of which are directed from X to Y we say that X is an *ancestor* of Y. If X is an ancestor of Y then Y is a *descendant* of X. The set of nodes connected to Y by a bidirected arc are called the *siblings* of Y. Lastly, a *collider* is a node where colliding arrowheads meet. Z in Figure 4a is a collider as are C, D, and E in Figure 3.

A *path* between X to Y is a sequence of edges, connecting the two vertices. A path may go either along or against the direction of the arrows. A *directed path* from X to Y is a path consisting only of arrows pointed towards Y. A *back-door path* from X to Y is a path begins with an arrow pointing to X and ends with an arrow pointing to Y. For example, in Figure 3, $C \leftarrow B \rightarrow E$, $C \rightarrow D \rightarrow E$, $C \leftarrow B \rightarrow D \rightarrow E$, and $C \rightarrow D \leftarrow B \rightarrow E$ are all paths between C and E. However, only $C \rightarrow D \rightarrow E$ is a directed path, and only $C \leftarrow D \rightarrow E$ and $C \leftarrow B \rightarrow D \rightarrow E$ are back-door paths. The significance of directed paths stems from the fact that they convey the flow of causality, while the significance of back-door paths stems from their association with confounding.

A graph is *acyclic* if it does not contain any cycles, a directed path that begins and ends with the same node. A graph is *cyclic* if it contains a cycle. A model in which the causal graph is acyclic is called *recursive* while models with cyclic graphs are called *non-recursive.*

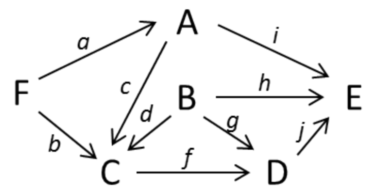

Figure 3: Model illustrating the rules of d-separation

(a) (b)

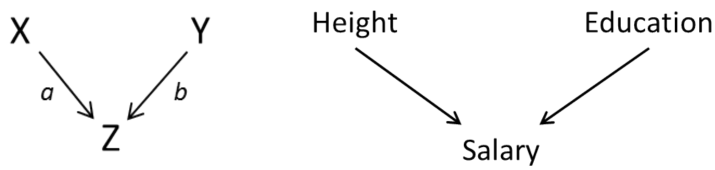

Figure 4: Examples illustrating conditioning on a collider

2.2 D-Separation

D-separation allows researchers to identify conditional independences implied by the model's structure from the causal graph and will be utilized extensively in the results to follow. The idea of d-separation is to associate "correlation" with "connectedness" in the graph, and independence with "separation".

Definition 1. *(Pearl, 2009, p. 16) A path p is said to be* d-separated *(or* blocked*) by a set of nodes Z if and only if*

(i) p contains a chain $i \rightarrow m \rightarrow j$ or a fork $i \leftarrow m \rightarrow j$ such that the middle node m is in Z, or

(ii) p contains a collider $i \rightarrow m \leftarrow j$ such that the middle node m is not in Z and such that no descendant of m is in Z.

A set Z is said to d-separate X from Y if and only if Z blocks every path from a node in X to a node in Y.

(i) captures the intuition that the correlation between two variables, X and Y, may vanish when conditioning on common causes or mediating variables. For example, the correlation between ice cream sale and drowning deaths is often used to show that correlation does not imply causation. When the weather gets warm people tend to both buy ice cream and play in the water, resulting in both increased ice cream sales and drowning deaths. This causal structure is depicted in Figure 5. Here, we see that Ice Cream Sales and Drownings are d-separated given either Temperature or Water Activities. As a result, if we only consider days with the same temperature and/or the same number of people engaging in water activities then the correlation between Ice Cream Sales and Drownings will vanish.

(ii) is due to the fact that conditioning on a collider or its descendant opens the flow of information between the parents of the collider. For example, X and Y are

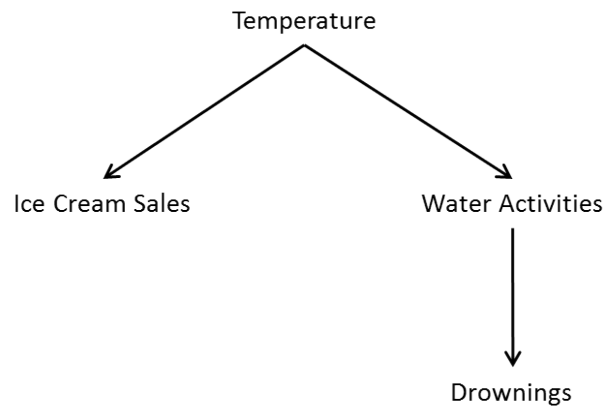

Figure 5: Graph illustrating why Ice Cream Sales and Drowning are uncorrelated given Temperature and/or Water Activities

uncorrelated in Figure 4a. However, conditioning on the collider, Z, correlates X and Y giving $X \not\perp\!\!\!\perp Y|Z$. This phenomenon is known Berkson's paradox or "explaining away". To illustrate, consider the example depicted in Figure 4b. It is well known that higher education often affords one a greater salary. Additionally, studies have shown that height also has a positive impact on one's salary. Let us assume that there are no other determinants of salary and that Height and Education are uncorrelated. If we observe an individual with a high salary that is also short, our belief that the individual is highly educated increases. As a result, we see that observing Salary correlates Education and Height. Similarly, observing an effect or indicator of salary, say the individual's Ferrari, also correlates Education and Height.

Berkson's paradox implies that paths containing colliders can be unblocked by conditioning on colliders or their descendants. Let π' be a path from X to Y that traces a collider. If for each collider on the path π', either the collider or a descendant of the collider is in the conditioning set Z then π' is unblocked given Z. The exception to this rule is if Z also contains a non-collider along the path π' in which case X and Y are still blocked given Z. For example, in Figure 3, the path $F \rightarrow C \leftarrow A \rightarrow E$ is unblocked given C or D. However, it is blocked given $\{A, C\}$ or $\{A, D\}$.

The following theorem makes explicit the relationship between conditional independence and d-separation.

Theorem 1. *Let G be the causal graph for a structural model over a set of variables V. If $X \in V$ and $Y \in V$ are d-separated given a set $Z \subset V$ in G then the model implies that X and Y are independent given Z.*

If X and Y are d-connected given Z then the set of points in the space of parameter values for which X and Y are uncorrelated given Z has Lebesgue measure zero (Spirtes et al., 1993). In other words, if X and Y are d-connected given Z then, according to the model, X and Y are almost surely correlated given Z but may be uncorrelated given Z for particular parameterizations. For example, it is possible that the values of the coefficients are such that the unblocked paths between X and Y perfectly cancel one another. For the remainder of the paper, whenever we invoke the notion of "almost surely" it will be in the same sense as we do here: to exclude coincidental cancellations.

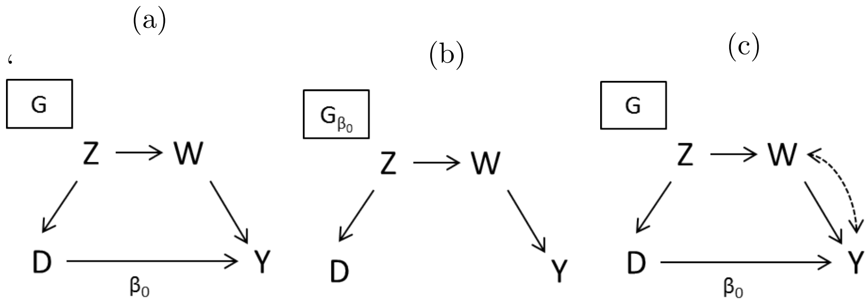

Figure 6: Graphs illustrating identification by the single-door criterion (a) β_0 is identified by adjusting for Z or W (b) The graph G_{β_0} used in the identification of β_0 (c) β_0 is identified by adjusting for Z (or Z and W) but not W alone

We use the graph depicted in Figure 3 as an example to illustrate the rules of d-separation. In this example, F is d-separated from E by A, B, and C. However, C is not d-separated from E by A and D since conditioning on D opens the collider $C \rightarrow D \leftarrow B$. Finally C is d-separated from E by conditioning on A, D, and B.

We conclude this section by noting that d-separation implies conditional independence in all recursive causal models, parametric or not (Pearl, 2009). In linear models, d-separation implies conditional independence in non-recursive models, as well as recursive models (Spirtes, 1995). Further, all vanishing partial correlations implied by a structural model can be obtained using d-separation (Pearl, 2009, ch. 1.2.3). Finally, in recursive models with independent error terms, these conditional independences represent all of the model's testable implications (Geiger and Pearl, 1993).

3 The Single-door Criterion

The single-door criterion is a necessary and sufficient graphical condition for the identification of structural coefficients using regression. When all observable drivers of Y are included in the regression specification then the single-door criterion graphically characterizes conditional exogeneity[4].

Theorem 2. *(Pearl, 2000) (Single-door Criterion) Let G be any graph for a linear model in which β_0 is the structural coefficient associated with link $D \rightarrow Y$, and let G_{β_0} denote the graph that results when the arrow from D to Y is deleted from G. The coefficient β_0 is identifiable if there exists a set of variables Z such that (i) Z contains no descendant of Y and (ii) Z d-separates D from Y in G_{β_0}. If Z satisfies these two conditions, then β_0 is equal to the regression coefficient $r_{YD.Z}$, and we say that Z is a* single-door admissible *with respect to β_0. Conversely, if Z does not satisfy these conditions, then $r_{YD.Z}$ is almost surely not a consistent estimand of β_0.*

Consider Figure 6a. As an observable driver of Y, W is generally included in the conditioning set when estimating β_0 using regression. The single-door criterion

[4]Unlike conditional exogeneity, however, the single-door criterion does not require that all observable drivers of Y be included in the regression. This allows us to obtain additional estimands for β_0, which will prove valuable for robustness testing.

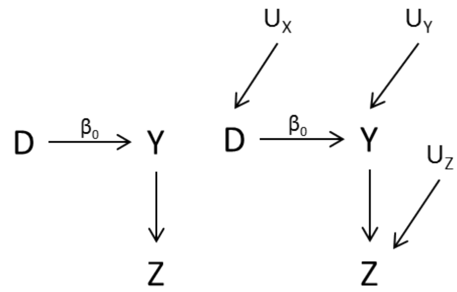

Figure 7: Example showing that adjusting for a descendant of Y induces bias in the estimation of β_0

confirms that, in this case, exogeneity holds and $r_{YD.W} = \beta_0$ since W blocks the spurious path $D \leftarrow Z \rightarrow W \rightarrow Y$ and D is d-separated from Y by W in Figure 6b. Theorem 2 tells us, however, that Z can also be used for adjustment since Z also d-separates D from Y in Figure 6b. Consider, however, Figure 6c. Z and $\{Z, W\}$ satisfy the single-door criterion but W does not. Being a collider, W unblocks the spurious path, $D \leftarrow Z \rightarrow W \leftrightarrow Y$, in violation of Theorem 2, leading to bias if adjusted for. In conclusion, β_0 is equal to $r_{YD.Z}$ and $r_{YD.WZ}$ in Figures 6a and 6c. However, β_0 is equal to $r_{YD.W}$ in Figure 6a only.

The intuition for the requirement that Z not be a descendant of Y is depicted in Figures 7a and 7b. We typically do not display the error terms, which can be understood as latent causes. In Figure 7b, we show the error terms explicitly. It should now be clear that Y is a collider and conditioning on Z will create spurious correlation between D, U_Y, and Y leading to bias if adjusted for.

4 Robustness Tests

In this section, we provide a criterion to discern when a shift in the coefficient r_0 is indicative of endogeneity bias and when it is not. Additionally, we demonstrate how and when robustness tests are able to detect omitted variables, including unobservable variables.

4.1 Shift-Producing Regressors

Once we have determined that a particular regression coefficient, $r_{YD.W_1}$, identifies β_0, we may wish to conduct robustness tests in order to check the model assumptions that imply $r_{YD.W_1} = \beta_0$. However, a shift in r_0 due to the addition of covariates is not necessarily indicative of model misspecification. In some cases, the model may imply that adding a particular covariate, W_2, will induce such a shift in r_0. As a result, the shift should be expected and is not indicative of endogeneity bias. We will call such covariates *shift-producing.*

Definition 2. *Given a hypothesized model M that implies $r_{YD.W_1} = \beta_0$, where β_0 is the parameter of interest, a covariate, W_2, is* shift-producing *for $r_{YD.W_1}$ if the model*

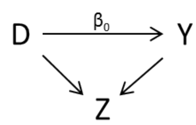

Figure 8: $r_{YD.Z} \neq \beta_0$ since Z is a collider

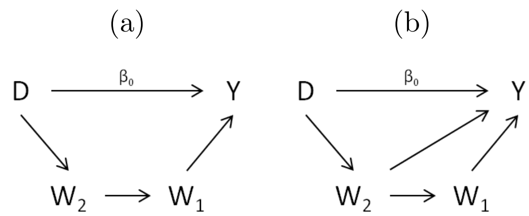

Figure 9: Descendants of D are necessary for both identification of β_0 and robustness testing

implies that adding it to the regression of Y on D and W_1 will shift the coefficient of D.

The next theorem uses the model graph to identify regressors that would be shift-producing for almost all parameterizations of the model.

Theorem 3. *Given a hypothesized structural model M that implies $r_{YD.W_1} = \beta_0$, W_2 is almost surely shift-producing for $r_{YD.W_1}$ if $W_1 \cup W_2$ does not satisfy the single-door criterion with respect to β_0. Conversely, a set of regressors, W_2, is not shift-producing if $W_1 \cup W_2$ satisfies the single-door criterion with respect to β_0.*

The single-door criterion gives a necessary and sufficient condition for identification using regression. Therefore, if $W_1 \cup W_2$ satisfies the single-door criterion then $r_{YD.W_1W_2} = \beta_0 = r_{YD.W_1}$ and adding W_2 to the regression of Y on D and W_1 should not shift r_0. Conversely, we know that defiance of the single-door criterion implies that $r_{YD.W_1W_2} \neq \beta_0 = r_{YD.W_1}$ (for almost all parameterizations of the model) and, therefore, adding W_2 will shift r_0 even when $r_{YD.W_1}$ gives an unbiased estimate of β_0.

Lu and White (2014) indeed identified some shift-producing regressors by noting that variables "driven by D" (i.e. descendants of D) should not be included in robustness tests for β_0. This precaution is often justified, but may be overly restrictive. Graphical analysis allows us to see when variables driven by D are problematic. Effects of D may be colliders, as in Figure 8 or descendants of Y, as in Figure 7a. In Figure 8, conditioning on the collider, Z, opens the spurious path, $D \to Z \leftarrow Y$. Therefore, Z is not single-door admissible and $r_{YD.Z} \neq \beta_0$. Similarly, Z is not single-door admissible in Figure7a since it is a descendant of Y. As we mentioned in Section 3, conditioning on Z violates exogeneity by inducing correlation between D and U_Y, since Z is a descendant of the collider, Y, as shown in Figure 7b. As a result, in both cases, $r_{YD} = \beta_0$ but $r_{YD.Z} \neq \beta_0$.

However, indiscriminately discounting descendants of D is a flawed strategy. In some cases, descendants of D are necessary for identification and can be useful for robustness testing. For example, if our model is the one depicted in Figure 9a, W_1 or W_2 must be included in the regression in order to obtain an unbiased estimate of β_0. In Figure 9b, W_2 must be added to the regression in order to obtain an unbiased

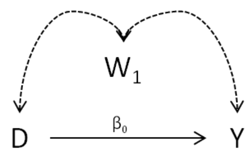

Figure 10: Z is not a descendant of D but adding Z to regression induces bias

estimate of β_0. Moreover, in the subsequent section on informative regressors, we will see that adding W_2 to the regression of Y on D and W_1 in Figure 9a constitutes an informative robustness test, demonstrating that descendants of D can also be used for robustness testing.

Finally, discounting variables driven by D is not sufficient to exclude all shift-producing regressors as non-descendants of D may also be shift-producing. For example, in Figure 10, $\beta_0 = r_{YD}$ but $\beta_0 \neq r_{YD.W_1}$ even though W_1 is not a descendant of D since conditioning on W_1 opens the spurious path, $D \leftrightarrow W_1 \leftrightarrow Y$. As a result, even if we suspect that W_1 may be a confounder, we cannot add it to the regression of Y on D to test whether this is the case.

In conclusion, while variables driven by D may be shift-producing, excluding all such variables for identification or robustness testing is both overly restrictive and insufficient for the exclusion of shift-producing regressors. Instead, Theorem 3 provides a simple criterion that is necessary and sufficient for a regressor to be shift-producing.

4.2 Informative Regressors

A robustness test is informative when a shift in the regression coefficient, r_0, is indicative of bias, and simultaneously, stability of the coefficient provides evidence that $r_0 = \beta_0$. Such tests are characterized by two properties:

(i) The hypothesized model must imply that r_0 equals β_0 before and after the addition of regressors.

(ii) There must be the possibility that the robustness test will fail in the case of misspecification[5].

We know from Theorem 2 that the first condition is satisfied if and only if both W_1 and $W_1 \cup W_2$ are single-door admissible. Additionally, if the model is incorrect and W_1 is not actually single-door admissible but $W_1 \cup W_2$ is, then the robustness test will fail. This is the case whenever the model implies that controlling for W_1 is enough to obtain an unbiased estimate of β_0 but in reality W_2 is necessary as well. For example, suppose that our model is the one depicted in Figure 9a but, in fact, we are missing an arrow from W_2 to Y, as in Figure 9b. Then comparing $r_{YD.W_1}$ and $r_{YD.W_1W_2}$ will detect this misspecification since $r_{YD.W_1} \neq r_{YD.W_1W_2}$[6]. In general, the second condition is also

[5]After identifying a set W_1 that satisfies the conditional exogeneity assumption, Lu and White (2014) advocate adding variables W_2 such that $D \perp\!\!\!\perp W_2 | W_1$ to the regression of Y on D and W_1 for robustness testing. While $D \perp\!\!\!\perp W_2 | W_1$ does imply that $D \perp\!\!\!\perp U | W_1, W_2$ when $D \perp\!\!\!\perp U | W_1$ (Dawid, 1979) so that the model implies $r_0 = \beta_0$ before and after the addition of W_2, the proposed robustness is non-informative. Pearl and Paz (2014) show that $D \perp\!\!\!\perp W_2 | W_1$ (see also Lemma 1 below) implies $r_{YD.W_1} = r_{YD.W_1W_2}$, even when conditional exogeneity is violated. As a result, the proposed robustness test will ALWAYS pass.

[6]This example also demonstrates our conclusion from the previous section. In some cases, descendants of D are not only necessary for identification of β_0 but also useful for robustness testing.

satisfied if both W_1 and $W_1 \cup W_2$ are single-door admissible in the hypothesized model, and we obtain the following theorem, which holds for almost all parameterizations of the data-generating model:

Theorem 4. *Given a hypothesized structural model M that implies $r_{YD.W_1} = \beta_0$, the robustness test comparing $r_{YD.W_1}$ and $r_{YD.W_1W_2}$ is informative for the hypothesis that $r_{YD.W_1} = \beta_0$ if and only if $W_1 \cup W_2$ is single-door admissible with respect to β_0.*

Proof. First, we show the sufficiency of Theorem 4. Property (i) of informative robustness tests is satisfied since both W_1 and $W_1 \cup W_2$ satisfy the single-door criterion in the hypothesized model. Property (ii) is satisfied since we can always construct a possibly "correct" model by adding arrows from W_2 to D and Y. If this model is the correct model, then the robustness test will fail.

The necessity of Theorem 4 holds almost surely since property (i) holds almost surely if and only if both W_1 and $W_1 \cup W_2$ are single-door admissible in the hypothesized model. □

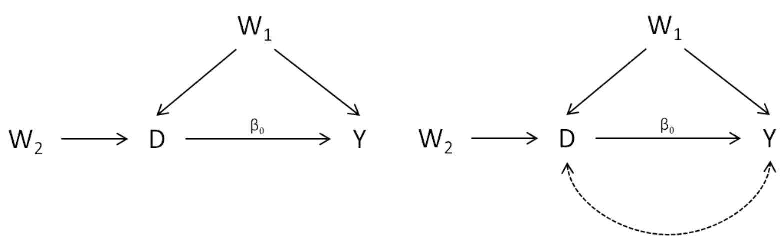

(a) (b)

Figure 11: Given the model depicted in (a), a confounding path not blocked by W_1 as in (b), can be detected by comparing $r_{YD.W_1}$ and $r_{YD.W_1W_2}$.

The above example, depicted by Figures 9a and 9b, demonstrates that informative robustness tests are able to detect misspecification when the added variable, originally thought to be irrelevant, is, in fact, an omitted variable or blocks the influence of an omitted variable. This should not be surprising since adding such a variable to the regression shifts r_0 by negating or reducing the omitted variable bias. As a result, stability of r_0 supports our hypothesis that W_2 is not an omitted variable and does not need to be added to the regression to estimate β_0. However, what about other omitted variables? Can the stability of r_0 tell us anything about the presence of omitted variables other than W_2 (or those blocked by W_2)? Surprisingly, the answer is yes.

Suppose that our model, given by Figure 11a, is misspecified, and the correct specification is given by Figure 11b. In this case (Figure 11b), conditioning on W_1 makes W_2 an instrumental variable. Since adding an instrumental variable to the regression of Y on D and W_1 in the presence of confounding amplifies the bias of r_0 (Bhattacharya and Vogt, 2007; Wooldridge, 2009; Pearl, 2010), the bidirected edge between D and Y ensures that $r_{YD.W_2} \neq r_{YD.W_1W_2}$ and comparing the regressions allows us to detect the omitted variable bias, even though W_2 is neither the omitted variable nor blocks its influence. Moreover, note that the robustness test was able to detect the omitted variable bias even though the variable is unobservable and cannot be added to the regression.

In general, if the omitted variable is a common cause of D and Y, and the added covariate, W_2, is a non-shift-producing cause (or proxy of a cause) of D, then detection is possible, even when W_2 neither is the omitted variable nor blocks its influence. The lemma below characterizes when $r_{YD.W_1} = r_{YD.W_1W_2}$ and will be used to show how and when omitted common causes can be detected using robustness tests.

Lemma 1. *For a variable S and a set of variables W, $r_{YD.W} = r_{YD.WS}$ if and only if*

(i) $\sigma_{SD.W} = 0$ *or*

(ii) $\sigma_{YS.DW} = 0$,

where $\sigma_{XY.Z}$ is the covariance of X and Y given Z[7]*.*

Proof. The sufficiency of Lemma 1 follows from Theorem 4 of (Pearl and Paz, 2014). We now show its necessity by showing that if $r_{YD.W} = r_{YD.WS}$, then either $\sigma_{SD.W} = 0$ or $\sigma_{YS.DW} = 0$.

In general, $r_{YD.WS} = \frac{r_{YD.W} - r_{YS.W}r_{SD.W}}{1 - r^2_{SD.W}\frac{\sigma^2_{D.W}}{\sigma^2_{S.W}}}$. Since we assume that $r_{YD.W} = r_{YD.WS}$ we get:

$$
\begin{aligned}
r_{YD.W} &= \frac{r_{YD.W} - r_{YS.W}r_{SD.W}}{1 - r^2_{SD.W}\frac{\sigma^2_{D.W}}{\sigma^2_{S.W}}} \\
0 &= -r_{YD.W} + r_{YD.W}r^2_{SD.W}\frac{\sigma^2_{D.W}}{\sigma^2_{S.W}} + r_{YD.W} - r_{YS.W}r_{SD.W} \\
&= r_{YD.W}r^2_{SD.W}\frac{\sigma^2_{D.W}}{\sigma^2_{S.W}} - r_{YS.W}r_{SD.W} \\
&= r_{SD.W}(r_{YD.W}r_{SD.W}\frac{\sigma^2_{D.W}}{\sigma^2_{S.W}} - r_{YS.W}) \\
&= r_{SD.W}(r_{YS.W} - r_{YD.W}r_{SD.W}\frac{\sigma^2_{D.W}}{\sigma^2_{S.W}})
\end{aligned}
$$

From the above, we see that either $r_{SD.W} = 0$, which implies that $\sigma_{SD.W} = 0$, or $r_{YS.W} - r_{YD.W}r_{SD.W}\frac{\sigma^2_{D.W}}{\sigma^2_{S.W}} = 0$. We also have:

$$
\begin{aligned}
& r_{YS.W} - r_{YD.W}r_{SD.W}\frac{\sigma^2_{D.W}}{\sigma^2_{S.W}} \\
&= r_{YS.W}\sigma^2_{S.W} - r_{YD.W}\sigma^2_{D.W}r_{SD.W} \\
&= \sigma_{YS.W} - \sigma_{YD.W}\frac{\sigma_{SD.W}}{\sigma^2_{D.W}} \\
&= \sigma_{YS.DW}
\end{aligned}
$$

[7] Lemma 1 implies that $r_{YD.W} = r_{YD.WS}$ can be characterized in terms of conditional independences. As a result, an informative robustness test simply checks that certain model-implied conditional independences are satisfied in the data. Since all conditional independences implied by the model can be identified using d-separation, robustness testing would be unnecessary if we check all of the testable implications identified using d-separation. (An efficient algorithm that utilizes the graph to enumerate a set, not necessarily minimal, of conditional independences that imply all others is given by Kang and Tian (2009).) However, some model-implied conditional independences may not be relevant to the conditional exogeneity assumption. Therefore, while robustness tests may fail to address certain testable implications that can be identified using d-separation they have the benefit of testing model assumptions that are used specifically in the identification of β_0. See (Pearl, 2004) and (Chen et al., 2014) for more on testing such assumptions.

Therefore, if $r_{YD.W} = r_{YD.WS}$ then either $\sigma_{SD.W} = 0$ or $\sigma_{YS.DW} = 0$. □

The theorem below formalizes the conditions under which a robustness test can detect an omitted common cause, even when the added variable is neither the omitted variable nor blocks its influence.

Theorem 5. *Given a hypothesized structural model M that implies $r_{YD.W_1} = \beta_0$, an informative robustness test comparing $r_{YD.W_1}$ and $r_{YD.W_1W_2}$ is almost surely able to detect omitted common bias of $r_{YD.W_1}$ AND $r_{YD.W_1W_2}$ if and only if there exists a path (in the correct model) from W_2 to D ending in an arrow or bidirected edge into D that is not blocked by W_1*[8]*.*

Proof. Omitted common cause bias of $r_{YD.W_1}$ and $r_{YD.W_1W_2}$ implies that model is misspecified so that there is at least one back-door path between D and Y that is not blocked by W_1 and $W_1 \cup \{W_2\}$. In this case, we have a model M, which is our misspecified model, and a model M', which is the correct model. M and M' are the same except M' contains at least one back-door path between D and Y that is not blocked by W_1 and $\{W_1, W_2\}$ that M does not.

We will show that $r_{YD.W_1} \neq r_{YD.W_1W_2}$ by showing that the two conditions of Lemma 1 are violated. First, there exists a path from W_2 to D that is not blocked by W_1 in M' by assumption and (i) is violated. Second, since there is a path from W_2 to D ending in an arrow or bidirected edge into D that is not blocked by W_1 and a back-door path between D and Y that is not blocked by W_1 or W_2 we know that W_2 is connected to Y given D and W_1 through a path $W_2 ... \to D \leftarrow ... \to Y$ and (ii) is violated. As a result, $r_{YD.W_1} \neq r_{YD.W_1W_2}$.

We now prove the necessity of Theorem 5. According to Lemma 1, if $r_{YD.W_1} \neq r_{YD.W_1W_2}$ then $\sigma_{W_2D.W_1} \neq 0$ and $\sigma_{W_2Y.DW_1} \neq 0$. $\sigma_{W_2D.W_1} \neq 0$ tells us that there exists a path from W_2 to D that is not blocked by W_1. We need to show that this path ends in an arrow or bidirected edge into D in M'.

$\sigma_{W_2Y.DW_1} \neq 0$ tells us that there is a path from W_2 to Y that is not blocked by D and W_1. Since W_2 is an informative regressor with respect to $r_{YD.W_1}$, W_1 and $W_1 \cup \{W_2\}$ are single-door admissible in M and we know that W_1 must block any path between W_2 and Y in M. This implies that W_1 blocks all paths between W_2 and Y in M' as well, since the only difference between M' and M are back-door paths between D and Y that are not blocked by W_1 and $W_1 \cup \{W_2\}$. Now, the only possibility of an unblocked path between W_2 and Y given D and W_1 is that there exists a path from W_2 to D ending in either an arrow or bidirected edge into D so that D is a collider, which opens a path $W_2 ... \to D \leftarrow ...Y$ that goes through an unaccounted for back-door path. □

[8] If the conditions of Theorem 5 (W_2 is connected to D via a path ending in an arrow or bidirected edge into D that is not blocked by W_1 and W_1 is single-door admissible) hold in the hypothesized model, then the model implies that W_2 is an instrumental variable given W_1 (Brito and Pearl, 2002; Chen and Pearl, 2014). As a result, the Hausman test for endogeneity, where the estimate of β_0 using the instrumental variable, $\frac{r_{YW_2.W_1}}{r_{DW_2.W_1}}$, is compared to the estimate using regression, $r_{YD.W_1}$, is applicable. In fact, both the Hausman and robustness tests check the same model-implied constraints. The Hausman test checks that $\frac{r_{YW_2.W_1}}{r_{DW_2.W_1}} = r_{YD.W_1}$. Rearranging terms we get $r_{YW_2.W_1} - r_{YD.W_1} r_{DW_2.W_1} = 0$, which is true if and only if $\sigma_{YW_2.DW_1} = 0$. Likewise, the robustness test checks that $r_{YD.W_1} = r_{YD.W_1W_2}$, which according to Lemma 1 also holds if and only if $\sigma_{YW_2.DW_1} = 0$. The same constraint could also be obtained using d-separation (see footnote 7).

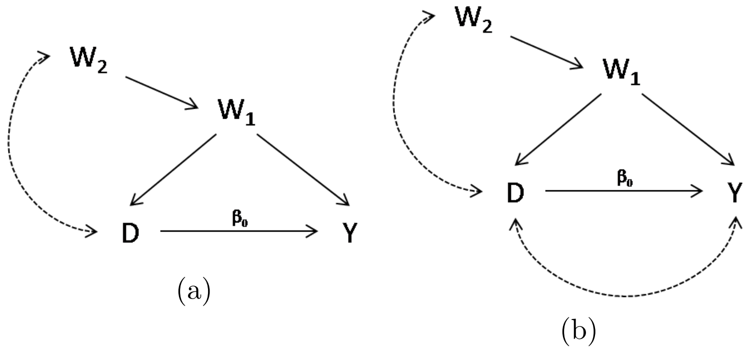

Figure 12: The bidirected edge from W_2 to D allows us to detect the latent confounder depicted in (b) from the hypothesized model depicted in (a) by adding W_2 to the regression of Y on D and W_1.

Returning to the hypothesized model depicted in Figure 11a, we see that adding W_2 to the regression of Y on D and W_1 is able to detect the latent confounder in the correct model, depicted in Figure 11b, since there is an arrow from W_2 to D. Similarly, suppose that we hypothesize the model depicted in Figure 12a. Both $\{W_1\}$ and $\{W_1, W_2\}$ are single-door admissible sets so adding W_2 to the regression of Y on D and W_1 represents an informative robustness test. Moreover, Theorem 5 tells us that stability of r_0 provides evidence that $r_{YD.W_1}$ does not suffer from omitted variable bias due to W_2 or other common causes since W_2 is connected to D by a bidirected edge. If, on the other hand, the correct model is the one depicted in Figure 12b, then adding W_2 to the regression will shift r_0 and detect the latent omitted variable.

Theorems 4 and 5 provide us with the means to answer the question: Have all necessary confounders been accounted for? If we believe that a set W_1 is adequate to cover all confounders, then we can test this assumption by adding non-shift-producing variables to the regression. If a variable, W_2, does not shift r_0, then we can rest assured that W_2 is not an omitted variable. If W_2 is a cause or proxy of a cause of D, then we further know that there are no other unaccounted for common causes.

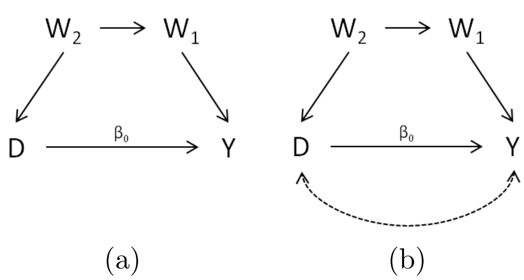

Figure 13: If the true model is given by (b) but the hypothesized model is given by (a) then using W_1 for identification of β_0 and adding W_2 to test robustness is able to detect misspecification but not the other way around.

Lemma 1 and Theorem 5 also demonstrate that some informative robustness tests may be "more informative" than others. For example, suppose that our hypothesized model is given by Figure 13a but the correct model is the one depicted in Figure 13b.

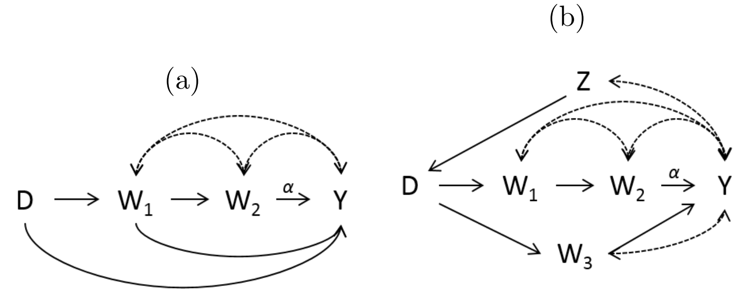

Figure 14: In both graphs, the total effect of D on Y is identifiable even though some of the individual coefficients comprising the effect are not. In (a), the total effect of D on Y is given by r_{YD} while in (b), the total effect of D on Y is given by $r_{YD.Z}$.

Using Theorem 5, we see that comparing $r_{YD.W_1}$ and $r_{YD.W_1W_2}$ allows us to detect the missing latent confounder since W_2 is a parent of D. However, while both $\{W_2\}$ and $\{W_1, W_2\}$ are single-door admissible in the hypothesized model, $r_{YD.W_2} = r_{YD.W_1W_2}$ whether there is a latent confounder or not, and we cannot detect a possible latent confounder by comparing these two regressions. This is easily verified using Lemma 1 since $W_1 \perp\!\!\!\perp D | W_2$.

In general, if both W_1 and W_2 are single-door admissible but W_2 blocks all paths between W_1 and D, adding W_2 to the regression of Y on D and W_1 may allow the detection of an unobservable confounder but adding W_1 to the regression of Y on D and W_2 will not. This is formalized in the corollary below.

Corollary 1. *Given a hypothesized structural model M, where W_1 and W_2 are both single-door admissible sets for β_0 but W_2 blocks all paths between W_1 and D, the robustness test comparing $r_{YD.W_1}$ and $r_{YD.W_1W_2}$ is more informative than the robustness test comparing $r_{YD.W_2}$ and $r_{YD.W_1W_2}$. The latter detect omitted variable bias only when W_2 is the omitted variable, while the former test may be able to detect bias due to other common causes.*

Proof. In this corollary, we assume that the only difference between the proposed model, M, and the correct model, M' is that M' has a back-door path between D and Y not blocked by W_1 or W_2. As a result, $D \perp\!\!\!\perp W_2 | W_1$ and $r_{YD.W_2} = r_{YD.W_1W_2}$. However, if there is an unblocked path between W_2 and D ending in an arrow or bidirected edge into D then Theorem 5 tells us that $r_{YD.W_1} \neq r_{YD.W_1W_2}$. □

Moreover, previous work by Kuroki and Miyakawa (2003) and Hahn (2004) have shown that, in this case, $\hat{r}_{YD.W_2}$ is preferable to $\hat{r}_{YD.W_1}$ for estimating β_0 in terms of asymptotic efficiency. The intuition here is that W_2 is "closer" to Y, hence more effective in reducing variations due to uncontrolled factors.

5 Total Effects

Thus far we have considered the case where we wish to evaluate our estimate of the direct effect of D on Y. However, in some cases, we may be interested in the total effect of D on Y (given by sums of products of coefficients along all directed paths from D to Y) rather than the direct effect (given by a single structural coefficient). In this section, we will discuss how the above results can be used for robustness tests when the quantity of interest is the total effect. First, we introduce the back-door criterion, which is a necessary and sufficient condition for the identifiability of a total effect using regression.

Theorem 6. *(Pearl, 2009, ch. 3) (Back-door Criterion) For any two variables D and Y in a model with causal diagram G, the total of effect of D on Y is identifiable by regression if and only if there exists a set of measurements Z such that*

(i) no member of Z is a descendant of D[9]*; and*

(ii) Z d-separates D from Y in the subgraph $G_{\underline{D}}$ formed by deleting from G all arrows emanating from D.

Moreover, if the two conditions are satisfied, then the total effect of D on Y is given by $r_{YD.Z}$ and we say that Z is a back-door admissible *set.*

For example, in Figure 14a, the total of effect of D on Y is r_{YD} since D is d-separated from Y when arrows leaving D are removed from the graph. Likewise, the total effect of D on Y in Figure 14b is $r_{YD.Z}$ since Z blocks all paths between D and Y when arrows emanating from D are removed.

The above examples also demonstrate that we may be able to identify the total effect even when the individual coefficients comprising it are not identified[10]. In both cases, the total effect was identifiable even though α is not.

Theorem 6 allows us to adapt Theorems 3 and 4 when the quantity of interest is a total rather than direct effect. We simply use the back-door criterion instead of the single-door. Similarly, Theorem 5 can be adapted and, in fact, strengthened since omitted variable bias when estimating total effects is due only to omitted common causes.

Theorem 7. *Given a hypothesized structural model M with β_0 the total effect of D on Y, an informative robustness test comparing $r_{YD.W_1}$ and $r_{YD.W_1W_2}$ is able to detect omitted variable bias of $r_{YD.W_1}$ AND $r_{YD.W_1W_2}$ if and only if there exists a path from W_2 to D ending in an arrow or bidirected edge into D that is not blocked by W_1*[11]*.*

Lastly, Corollary 1 also holds when β_0 is the total effect of D on Y. We simply replace the single-door criterion with the back-door criterion.

6 Example

In this section, we give two examples using simulated data to demonstrate how the above material can be used to avoid shift-producing regressors and conduct informative robustness tests. In the first example, we will demonstrate how Theorem 3 can aid in the detection of shift-producing regressors. In the second example, we will demonstrate

[9] Descendants of D can be included in the conditioning set, Z, so long as spurious paths between D and Y opened by descendants of D are blocked by Z. However, such variables can always be excluded without inducing bias.

[10] The ability to answer policy questions even when individual parameters are not identified was first noted by Marschak (1942). This principle was dubbed "Marschak's Maxim" by Heckman (2000).

[11] Again, the sufficiency of this theorem holds almost surely since the correct model may parameterized in such a way that certain paths cancel one another.

how Theorem 4, Theorem 5, and Corollary 1 can be used to ensure that a given estimate is free of omitted variable bias.

For each simulation, we utilize the tr_robust_test function from the *testrob* Matlab procedure given by Lu and White (2014). This function conducts a Hausman-style test to determine whether a given robustness test passes or fails. For each example, we conducted 100,000 trials with a sample size of 1,000. Error terms were drawn from Gaussian distributions with mean 0 and variance 1. Additionally, the coefficients for each equation/arrow are drawn uniformly at random from the interval $(0.5, 1)$. We give the percentage of trials for which a given robustness test passes, as well as whether the proposed regression coefficient accurately estimates β_0, according to a t-test at a 5% significance level.

6.1 Example 1

Suppose that we have the following hypothesized model, depicted graphically in Figure 10, and we wish to estimate β_0.

Model 5.

$$\begin{aligned} D &= U_D \\ W_1 &= U_1 \\ Y &= \beta_0 D + U_Y \\ \sigma_{U_D U_1} &\neq 0 \\ \sigma_{U_Y U_1} &\neq 0 \end{aligned}$$

First, we note that it is not immediately obvious from the equations whether we should condition on W_1 or not. W_1 is correlated with both D and Y, and it is not driven by D. In this sense, it seems like a confounder that should be added to the regression. However, the single-door criterion tells us that $r_{YD.W_1} \neq \beta_0$ since conditioning on W_1 opens the spurious path, $D \leftrightarrow W_1 \leftrightarrow Y$. Instead, $\beta_0 = r_{YD}$. Indeed, when using data generated according to Model 5, we find that $r_{YD} = \beta_0$ in 94.98% of trials.

Now, suppose that we suspect that W_1 is a driver of Y and that the equation for Y is actually

$$Y = \beta_0 D + \alpha W_1 + U_Y.$$

In this case, $\beta_0 \neq r_{YD}$. As a result, we would like to add W_1 to the regression to test whether this is the case. However, Theorem 3 tells us that W_1 is shift-producing and that even if Model 5 is correct, the regression coefficient will shift. As a result, such a robustness test is non-informative. Indeed, adding W_1 to the regression of Y on D shifts r_0 in 100% of trials using data generated according to Model 5.

6.2 Example 2

In this next example, we will demonstrate how to conduct informative robustness tests to ensure the lack of omitted variable bias when estimating β_0. Suppose that we have the following hypothesized model, depicted in Figure 15a:

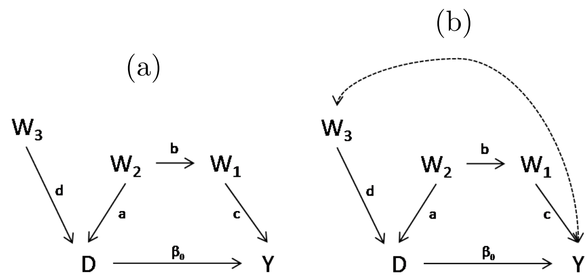

Figure 15: (a) Initial hypothesized graph used in Section 6.2 (b) Actual graph used in Section 6.2

Model 6.

$$
\begin{aligned}
W_1 &= bW_2 + U_1 \\
W_2 &= U_2 \\
W_3 &= U_3 \\
D &= aW_1 + dW_3 + U_D \\
Y &= \beta_0 D + cW_1 + U_Y
\end{aligned}
$$

In the hypothesized model, all error terms are independent of one another. However, let us suppose that, in reality, U_3 and U_Y are correlated, as shown in Figure 15b.

Using the single-door criterion and the graph for the hypothesized model, depicted in Figure 15a, we see that Model 5 implies that $r_{YD.W_1} = r_{YD.W_2} = r_{YD.W_1W_2} = \beta_0$. As a result, adding W_1 to the regression of Y on D and W_2 and adding W_2 to the regression of Y on D and W_1 both represent informative robustness tests. Corollary 1, however, tells us that the latter is "more informative" than the former. Indeed, the latter test detects the bias due to the omitted variable, W_3, in 80.40% of trials while the former does so in only 5.07% of trials.

After observing a shift when adding W_2 to the regression of Y on D and W_1, we may very well wonder if the detected omitted variable bias of $r_{YD.W_1}$ is remedied by adding W_2 to the regression or if other variables are necessary, as well. In other words, we may wonder if the correct model is the one in which the equation for Y is

$$Y = \beta_0 D + cW_1 + eW_2 + U_Y$$

and/or U_3 is correlated with U_Y (depicted in Figure 16a). In each case, $\beta_0 = r_{YD.W_1W_2}$. To test this hypothesis, we consult Figure 16a, and we see that adding W_3 to the regression of Y on D, W_1, and W_2 is an informative robustness test that is capable of identifying omitted variable bias of $r_{YD.W_1W_2}$ due to either W_3 or another variable entirely.

Adding W_3 shifts r_0 in 100% of trials, and we conclude that $r_{YD.W_1W_2} \neq \beta_0$. At this point, we are out of variables to add to the regression. However, recall that adding W_2 to the regression of Y on D and W_1 shifts r_0 when $r_{YD.W_1}$ suffers from omitted variable bias due to W_2 OR other variables. As a result, it is possible that W_3, not W_2, was responsible for the omitted variable bias that shifted r_0 when comparing $r_{YD.W_1}$ and $r_{YD.W_1W_2}$. This hypothesized model is depicted in Figure 16b and $r_{YD.W_1W_3} = \beta_0$

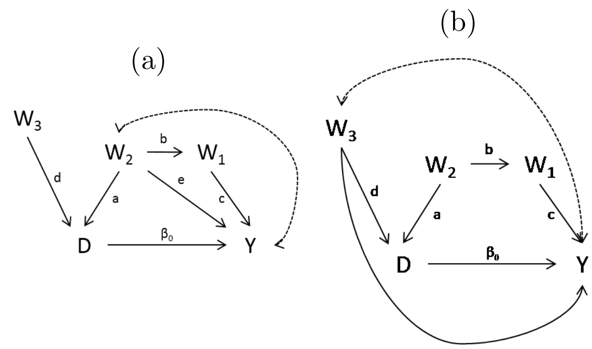

Figure 16: (a) Hypothesized graph where $r_{YD.W_1W_2} = \beta_0$ (b) Hypothesized graph where $r_{YD.W_1W_3} = \beta_0$

can be tested by adding W_2 to the regression. In 94.99% of trials, r_0 does not shift and the robustness test passes. As a result, we conclude that W_2 is not an omitted variable. Furthermore, we correctly conclude that there are no other omitted common causes because Theorem 5 tells us that the edge, $W_2 \rightarrow D$, ensures that our robustness test would have detected such bias. As expected, in 95.11% of trials $r_{YD.W_1W_3} = \beta_0$.

7 Conclusion

Robustness tests are a valuable tool in testing the structural validity of regression coefficients. However, the covariates involved in such tests must be carefully selected according to the economic model. Certain covariates will produce a shift in r_0 when added to the regression, even when it was structurally valid prior to the addition of regressors. A robustness test is informative only when the economic model implies that the regression coefficient is equal to the structural coefficient before and after the addition of regressors. We have given simple graphical criteria that allow researchers to quickly determine which regression coefficients are structurally valid, according to the model. As a result, they not only aid researchers in the selection of covariates for identification but also for informative robustness tests. Additionally, we have shown that robustness tests are not only able to detect bias due to omitted observable variables but unobservable variables as well. We provided a theorem characterizing when detection of omitted common causes is possible. Finally, we gave two extended examples using simulated data demonstrating how these results can be used to avoid shift-producing regressors and conduct informative robustness tests.

8 Acknowledgements

We would like thank Angus Deaton, James Heckman, Michael Keane, and Edward Leamer for helpful discussions. This research was supported in parts by grants from NSF #IIS-1302448 and ONR #N00014-10-1-0933 and #N00014-13-1-0153.

References

Bhattacharya, J. and Vogt, W. (2007). Do instrumental variables belong in propensity scores? Tech. Rep. NBER Technical Working Paper 343, National Bureau of Economic Research, MA.

Brito, C. and Pearl, J. (2002). Generalized instrumental variables. In *Uncertainty in Artificial Intelligence, Proceedings of the Eighteenth Conference* (A. Darwiche and N. Friedman, eds.). Morgan Kaufmann, San Francisco, 85–93.

Chen, B. and Pearl, J. (2014). Graphical tools for linear structural equation modeling. Tech. Rep. R-432, <http://ftp.cs.ucla.edu/pub/stat_ser/r432.pdf>, Department of Computer Science, University of California, Los Angeles, CA. Forthcoming, Psychometrika.

Chen, B., Tian, J. and Pearl, J. (2014). Testable implications of linear structual equation models. R-428, <http://ftp.cs.ucla.edu/pub/stat_ser/r428.pdf>, CA. To appear in Proceedings of AAAI-2014.

Dawid, A. (1979). Conditional independence in statistical theory. *Journal of the Royal Statistical Society, Series B* **41** 1–31.

Geiger, D. and Pearl, J. (1993). Logical and algorithmic properties of conditional independence and graphical models. *The Annals of Statistics* 2001–2021.

Hahn, J. (2004). Functional restriction and efficiency in causal inference. *Review of Economics and Statistics* **86** 73–76.

Heckman, J. (2000). Causal parameters and policy analysis in economics: A twentieth century retrospective. *The Quarterly Journal of Economics* **115** 45–97.

Hoover, K. (2004). Lost causes. *Journal of the History of Economic Thought* **26** 149–164.

Hoover, K. D. and Phiromswad, P. (2013). Selecting instrumental variables: a graph-theoretic approach. *Available at SSRN* .

Kang, C. and Tian, J. (2009). Markov properties for linear causal models with correlated errors. *The Journal of Machine Learning Research* **10** 41–70.

Kuroki, M. and Miyakawa, M. (2003). Covariate selection for estimating the causal effect of control plans by using causal diagrams. *Journal of the Japanese Royal Statistical Society, Series B* **65** 209–222.

Lu, X. and White, H. (2014). Robustness checks and robustness tests in applied economics. *Journal of Econometrics* **178** 194–206.

Marschak, J. (1942). Studies in mathematical economics and econometrics – in memory of henry schultz. chap. Economic interdependence and statistical analysis. University of Chicago Press, Chicago, 135–150. Reprinted in D.F. Hendry and M.S. Morgan (Eds.), *The Foundations of Econometric Analysis*, Cambridge University Press, 427–439, 1995.

Orcutt, G. (1952). Toward a partial redirection of econometrics. *Review of Economics and Statistics* **34** 195–213.

Oster, E. (2013). Unobservable selection and coefficient stability: Theory and validation. Tech. rep., National Bureau of Economic Research.

Pearl, J. (2000). *Causality: Models, Reasoning, and Inference*. Cambridge University Press, New York. 2nd edition, 2009.

Pearl, J. (2004). Robustness of causal claims. In *Proceedings of the Twentieth Conference Uncertainty in Artificial Intelligence* (M. Chickering and J. Halpern, eds.). AUAI Press, Arlington, VA, 446–453.

Pearl, J. (2009). *Causality: Models, Reasoning, and Inference*. 2nd ed. Cambridge University Press, New York.

Pearl, J. (2010). On a class of bias-amplifying variables that endanger effect estimates. In *Proceedings of the Twenty-Sixth Conference on Uncertainty in Artificial Intelligence*. AUAI, Corvallis, OR, 425–432. <http://ftp.cs.ucla.edu/pub/stat_ser/r356.pdf>.

Pearl, J. and Paz, A. (2014). Confounding equivalence in causal inference. *Journal of Causal Inference J. Causal Infer.* **2** 75–93.

Spirtes, P. (1995). Directed cyclic graphical representation of feedback. In *Proceedings of the Eleventh Conference on Uncertainty in Artificial Intelligence* (P.Besnard and S. Hanks, eds.). Morgan Kaufmann, San Mateo, CA, 491–498.

Spirtes, P., Glymour, C. and Scheines, R. (1993). *Causation, Prediction, and Search*. Springer-Verlag, New York.

Tinbergen, J. (1939). *Statistical testing of business-cycle theories: Business cycles in the United States of America, 1919-1932*. Statistical Testing of Business-cycle Theories, II.: Business Cycles in the United States of America, 1919-1932, League of nations, Economic intelligence service.
URL `http://books.google.com/books?id=2U03AAAAIAAJ`

Walter, S. and Tiemeier, H. (2009). Variable selection: current practice in epidemiological studies. *European journal of epidemiology* **24** 733–736.

White, H. and Chalak, K. (2009). Settable systems: An extension of pearl's causal model with optimization, equilibrium and learning. *Journal of Machine Learning Research* **10** 1759–1799.

White, H. and Chalak, K. (2010). Testing a conditional form of exogeneity. *Economics Letters* **109** 88–90.

White, H. and Lu, X. (2011). Causal diagrams for treatment effect estimation with application to efficient covariate selection. *Review of Economics and Statistics* doi:10.1162/REST_a_00153, Forthcoming.

Wold, H. (1954). Causality and econometrics. *Econometrica: Journal of the Econometric Society* 162–177.

Wooldridge, J. (2009). Should instrumental variables be used as matching variables? Tech. Rep. <https://www.msu.edu/~ec/faculty/wooldridge/current%20research/treat1r6.pdf>, Michigan State University, MI.