JOURNAL OF ECONOMIC METHODOLOGY
2022, VOL. 29, NO. 4, 326–334
https://doi.org/10.1080/1350178X.2022.2088085

Routledge
Taylor & Francis Group

Pearl before economists: the book of why and empirical economics

Nick Huntington-Klein

Department of Economics, Seattle University, Seattle, WA, USA

ABSTRACT
Structural Causal Modeling (SCM) is an approach to causal inference closely associated with Judea Pearl and given an accessible instroduction in [Pearl, J., & Mackenzie, D. (2018). *The book of why: The new science of cause and effect*. Basic Books]. It is highly popular outside of economics, but has seen relatively little application within it. This paper briefly introduces the main concepts of SCM through the lens of whether applied economists are likely to find marginal benefit in these methods beyond standard economic approaches to causal inference. The most promising areas are those where SCM's causal diagrams alone offer significant value: covariate selection, the development of placebo tests, causal discovery, and identification in complex models.

KEYWORDS
Causality; potential outcomes; pearl; causal diagrams; do-calculus

JEL CODE
C10

1. Introduction

Empirical applications in economics are deeply concerned with the identification of causal relationships. By one measure, discussion of identification occurs in as many as half of new National Bureau of Economic Research (NBER) working papers (Currie et al., 2020).

Despite heavy interest in using data to infer causal relationships, or causal inference, there is a popular and growing approach to causal inference that is largely unknown within the field. This is the Structural Causal Model (SCM) approach, tightly associated with its originator, Judea Pearl. *The Book of Why* (Pearl & Mackenzie, 2018) is one of the more accessible introductions to SCM, at least among those that count Pearl as an author. Morgan and Winship (2015) offer another introduction that is focused on the applied social sciences and is also accessible, while aiming at an audience of researchers.

SCM relies on a diagram and a set of structural equations describing causal relationships, which can then be evaluated using the *do*-calculus, a set of repeatable operations that can be used to turn a statement about causal intervention into a statement about conditional relationships that can be calculated in observed data.

SCM sets itself apart from the potential-outcomes model (PO) associated with Donald Rubin (1974), which is the theoretical approach to causal inference most familiar to economists. One way of viewing the distinction is that PO focuses first on modeling counterfactuals, while SCM focuses first on modeling the causal relationships between variables. PO might, for example, envision a causal effect as comparing an observed outcome against a missing-value counterfactual to be imputed, while SCM might envision a causal effect as the outcome of a structural model that is modified to set some treatment variable to a counterfactual value (Pearl & Mackenzie, 2018, pp. 269–286).

CONTACT Nick Huntington-Klein nhuntington-klein@seattleu.edu

SCM formalizes the process of causal calculation, and is similar in many ways to the economic tradition of structural modeling. Both SCM and structural modeling in economics descend directly from Wright (1921) and use structural equations. Economists doing structural equation modeling, however, would not typically use Pearl's diagram analysis or the *do*-calculus.

Pearl's contributions to SCM since the 1990s have found immense popularity in a fairly short period of time, especially in fields that did not already have strong traditions of causal inference. Economics, on the other hand, is already familiar with one approach to causal inference, and any economist should naturally wonder at the marginal rather than absolute value of the new contribution.[1]

This paper is neither a thorough treatment of SCM nor an overall appraisal of its value. Instead, it introduces some of the theory and tools of SCM as they are shown in The Book of Why, and discusses how these might be of interest or use to applied economists who are already familiar with causal inference as economists have understood it for the past few decades, i.e. without direct reliance on Pearl's SCM. As opposed to other papers tackling similar questions (Heckman & Pinto, 2022; Imbens, 2020), the approach taken here is focused on the view of an applied econometrician, especially one who is looking to pick and choose contextually useful tools to aid in causal inference research. Any suggestion that an applied economist might not find Pearl's work useful for a particular problem does not imply a deficiency in SCM, but rather that this particular kind of researcher is unlikely to find their capabilities greatly expanded beyond where they already are by incorporating SCM for that problem.[2]

2. Central contributions

In this section I describe some of Pearl's contributions, described in *The Book of Why*.

2.1. Ladder of causation

Pearl characterizes questions about the world as falling along one of several rungs on the 'Ladder of Causation' (Pearl & Mackenzie, 2018, Ch. 1). The first of these rungs is 'association', which has to do with seeing and observing, and is similar to what economists might call 'descriptive work.'

The interesting deviation from standard economic thinking on causality is the distinction between the second and third rungs, which are 'intervention' and 'counterfactuals,' respectively. Intervention has to do with questions about what would occur on intervening to change the value of a variable. Many policy-relevant causal questions like 'if we raise the minimum wage, what will happen to employment?' lie on this rung. Pearl distinguishes this from 'counterfactuals' which ask what would have happened in a counterfactual world, for example 'if the minimum wage had not been raised, would employment have been higher?'

Economists might not be used to distinguishing these two types of causal questions and may see them as equivalent, but Pearl's work does distinguish them, providing an analytical framework for analyzing both. It also provides a philosophical underpinning for asking questions about causes that cannot be manipulated, like race, which make sense in rung 3 as an example of a 'parallel world' counterfactual, but not rung 2, since you cannot intervene to change someone's race. These immutable causes are commonly applied but are on shaky philosophical ground in the standard economic understanding of causality, while being supported in Pearl's system.

2.2. Causal diagrams

Possibly the most visible feature of Pearl's work is his use of diagrams, as described in *The Book of Why* Chapter 3, to represent structural causal relationships. Pearl's diagrams are ways of representing sets of structural equations, except that they are fully nonparametric, and use arrows instead of equal signs to emphasize the causal direction of structural relationships.

For example, consider the set of linear structural equations, assuming that the term on the left is causally determined by the expression on the right.

$$X = \gamma_0 + \gamma_1 Z + \nu \tag{1}$$

$$Y = \beta_0 + \beta_1 X + \beta_2 X \times W + \beta_3 Z + \beta_4 W + \varepsilon \tag{2}$$

This system could be represented as the diagram in Figure 1. This same diagram represents any set of structural equations in which the same set of causal relationships hold, in any parametric arrangement.

Compared to a set of structural equations, a causal diagram abstracts away functional form and parametric restrictions (note that the interaction between X and W in the structural equation is not shown expressly in the diagram, but the diagram is not inconsistent with it), and makes more explicit the direction of causal relationships in the system.

Causal diagrams can be combined with the structural equations they represent. Skipping that step and relying on causal diagrams alone leaves some important causal questions unanswered, as I will discuss in the 'Pearl and Economic Research' section.

One major benefit of causal diagrams is that they allow researchers to focus directly on questions of causal identification, abstracting away all other issues. This makes tasks such as covariate selection and the evaluation of identifying assumptions much more straightforward.

Another major benefit is pedagogical. Causal diagrams, since they do not rely on complicated sets of equations, can be much easier to explain in an econometrics classroom than structural equations or even PO applications. As a result, there are several economics-focused causal inference textbooks that make use of causal diagrams as a primary teaching tool (Cunningham, 2021; Huntington-Klein, 2021).

2.3. The do-Calculus

The *do*-calculus is a set of three simple rules that allow researchers to calculate the proper causal estimand given a causal diagram. This process begins by defining the *do*() operator, which is an explicitly causal mathematical object. $P(Y \mid do(X))$ describes the distribution of Y conditional on intervening to *set* X to a particular value, which is distinct from the $P(Y \mid X)$ we observe in data. In terms of the causal diagram, specifying $do(X)$ in effect deletes any arrows pointing towards X, since we have set X's value by intervention, rather than having it determined by the normal data generating process.

$P(Y \mid do(X))$ might describe the target of interest, 'the effect of X on Y', and the *do*-calculus is a series of operations that describes how to turn this counterfactual term into an estimand that can be calculated using observed values. The three operations performed in the *do*-calculus are:

(1) Delete any arrows heading into variables in the *do*() operator. Then, if a variable has no conditional relationship with the outcome, it can be ignored. $P(Y \mid do(X), A, B) = P(Y \mid do(X), A)$ if a diagram deleting arrows into X implies that $Y \perp\!\!\!\perp B \mid X, A$. In other words, we can ignore variables irrelevant to identification.
(2) Delete any arrows coming out of variables in a *do*() operator. Then, if that variable has no conditional relationship with the outcome, we can eliminate the *do*() operator. $P(Y \mid do(X), A) = P(Y \mid X, A)$ if a diagram deleting arrows coming out of X implies $Y \perp\!\!\!\perp X \mid A$. In

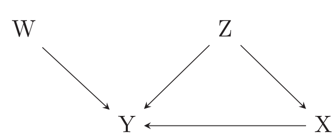

Figure 1. Example causal diagram.

other words, we can turn causal statements about X into observational statements if we have a set of covariates A that account for the non-causal relationships between X and Y.

(3) (3) If there is no pathway from the variable in the *do*() operator to the outcome in which all arrows point towares the outcome, that *do*() operator can be ignored. $P(Y \mid do(X)) = P(Y)$ if there is no direct causal path from X to Y. In other words, we can ignore an intervention X if changing the value of X does not change the value of Y.

If the *do*-calculus cannot produce an estimand calculable using observed values, then it cannot be done; the target is unidentifiable.[3]

The obvious benefit of the *do*-calculus is that a causal estimand can be derived purely from a causal diagram. This can be done even in unusual cases or with complex diagrams. This eases identification in cases where recognizable causal structures do not occur (i.e. conditioning on observables is not sufficient, nor is instrumental variables or other standard quasiexperimental designs). Throughout *The Book of Why*, Pearl presents examples of diagrams where identification is very difficult without application of the *do*-calculus.

3. Pearl and economic research

Pearl's SCM methods offer a comprehensive approach to structural causal identification that has found proponents in a wide array of fields. These methods have not been nearly as popular in economics. Speculatively, this may be because economists are already familiar with another set of approaches to causal inference, especially although not exclusively based around potential outcomes.

What is the *marginal* benefit of using Pearl's methods for researchers familiar with potential outcomes, or economic approaches to structural modeling? In a technical sense, the marginal benefit is zero, since potential outcomes and SCM are logically equivalent. Any theorem provable in one system is also provable in the other (Galles & Pearl, 1998).

However, even though these systems are logically equivalent, they are still different in the way they are used and what practices they encourage. They also differ in terms of which sorts of causal inference problems they make easy and which they make difficult.

In this section I discuss the kinds of problems that are easier in one approach or the other, and how that intersects with the kinds of problems economists routinely face.

3.1. Covariate selection

Probably the most obvious marginal benefit of the Pearl approach is in the realm of covariate selection. Applied econometricians have traditionally had little formal guidance when it comes to the selection of covariates. Causal diagrams, on the other hand, make the selection of covariates an automatic and well-informed process. There is real value to economists in using diagrams to improve and add structure to the covariate selection process, as White and Lu (2011) pointed out more than a decade ago.

As described in Chapter 3 of *The Book of Why*, a causal diagram, even without a set of accompanying structural equations, is sufficient to determine which variables the researcher must condition on to identify a causal effect ('achieve *d*-separation' in Pearl's terms). The process can even be carried out automatically (Textor et al., 2011).

Outside of the automatic selection of covariates, causal diagrams make the reasoning behind covariate selection much more clear. Diagrams make explicit the associational pathways that generate omitted variable bias, demonstrating why variables on those pathways must be controlled. Diagrams also make clear which controls are unnecessary (if they do not fall along a pathway that leads to bias) or can be harmful.

Economists already have a sense that variables caused by treatment are bad controls, but as demonstrated in Cinelli et al. (2020), causal diagrams greatly expand the detection of bad controls. The concept of collider variables is a prominent example (*The Book of Why* Chapter 6).

On a causal diagram, a variable is a collider on a pathway if variables on either side of it in the pathway cause it (the arrows 'collide' at that variable), for example the path $X \rightarrow Z \leftarrow Y$ in Figure 2. When this happens, the pathway is 'pre-closed' and does not bias the $X \rightarrow Y$ effect. Further, controlling for Z makes things worse by opening the path back up and introducing bias into the estimate.

Economists are familiar with selection bias, which is a form of collider bias. For example, in a randomized controlled trial study of whether personal tutoring in secondary school (X in Figure 2) improves intelligence (Y), the unconditional relationship between X and Y would provide an unbiased estimate of the effect. However, if the sample is then limited only to students who get into college (conditioning on Z), the estimate will become biased. Collider bias is not just limited to selection bias, however, and causal diagrams allow economists to recognize other applications.

3.2. Placebo tests

Placebo tests are when the researcher uses their intended method in such a way that it should produce an effect estimate of zero. For example, users of difference-in-difference methodology are familiar with the concept of looking for treatment effects in periods where the treatment did not occur; finding a nonzero effect is evidence that an assumption has been violated.

Placebo tests are a powerful way of checking the plausibility of untestable assumptions, like parallel trends in difference-in-differences. Causal diagrams reveal the full set of nonparametric placebo tests that can be implemented.

Any given causal diagram implies a set of conditional independences. For any two variables A and B on a diagram, there are only so many paths between those variables. Conditioning for a set of variables that closes all pathways should produce a conditional relationship between A and B of zero. If the relationship is nonzero, then one of the assumptions used to construct the diagram is incorrect.

For example, consider Figure 3, which simply replicates Figure 1. In this diagram, the only pathways between W and Z are $W \rightarrow Y \leftarrow Z$ and $W \rightarrow Y \leftarrow X \leftarrow Z$. Both of these pathways contain a collider and so should not induce any relationship between W and Z. The diagram implies that W is independent of Z. If the unadjusted relationship between W and Z is nonzero in the data, then some assumption is incorrect.

The causal diagram itself suggests a full set of these conditional dependencies in the data, which provides a host of placebo tests the researcher can apply, as long as all the relevant variables are observed.

3.3. Promising new methods

Pearl's methods naturally imply several additional causal identification strategies that economists may not be familiar with, which may allow for causal identification in new settings.

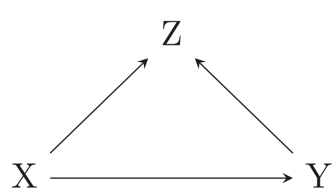

Figure 2. Causal Diagram where One Path Has a Collider.

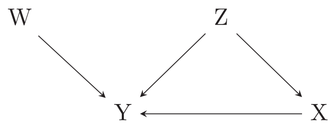

Figure 3. Example causal diagram.

The first of these is a research design, the front-door method (*The Book of Why* Chapter 5). The front-door method applies in cases like Figure 4 where the effect of interest (the effect of X on Y) is fully mediated by some observable variable(s) W. If Z is unobservable, then $X \rightarrow Y$ is not directly identifiable, but $X \rightarrow W$ is identified by the unadjusted relationship between X and W, and $W \rightarrow Y$ is identified by adjusting for X. Then, $X \rightarrow W$ is combined with $W \rightarrow Y$ to estimate the effect of X on Y. The front-door method is a good example of a design that is possible to derive with either potential outcomes or using a causal diagram and the *do*-calculus, but far easier in the latter (*The Book of Why* Chapter 7).

Actual applications of this method have been sparse. However, it is possible that this is simply because applied researchers are not attuned to looking for cases where the front-door method applies in the way that we look for instrumental variables or settings where quasiexperimental designs apply. Bellemare et al. (2020) offers one in-practice application of the front-door method to estimate the effect of the availability of trip-sharing on tipping activity in Uber rides.

Another area where causal diagrams imply a whole new way to generate theoretical results is in the area of causal discovery (see, e.g. Spirtes et al., 2000). Because, as previously discussed, causal diagrams make the process of covariate selection automatic and are capable of generating a set of implied conditional independences for any diagram, it becomes much easier to show that a given data set is inconsistent with a certain model.

Causal discovery algorithms narrow down the set of possible causal diagrams consistent with a data set. They use placebo tests to derive causal relationships directly from data without theory beyond selecting a list of variables. This is a powerful concept.

For example, consider a setting where an economist thinks that A, B, and C are the only relevant variables in the system. The data reveals that A and B are unrelated, but are related conditional on C, A and C are related whether conditional on B or not, and B and C are related whether conditional on B or not. This set of observed results is only consistent with a diagram where $A \rightarrow C$, $B \rightarrow C$, and there is no arrow between A and B.

Of course, in many cases the full model cannot be specified. The method also requires special ways of handling unobserved variables, as well as a definition of 'unrelated.' But this still offers an interesting alternative way that economists could derive insights that is not currently being used.

Finally, the *do*-calculus itself may unlock structural research in areas that were previously too difficult. There are many cases in economics research where the behavior being studied is the result of multiple interlocking decisions that causally relate in complex ways. Understanding identification in these models can be quite difficult. However, the *do*-calculus can be applied to derive causal estimands in fairly complex scenarios, including cases in which there are multiple treatments, or one treatment causes another, or there are long chains of causal relationships. *The Book of Why*

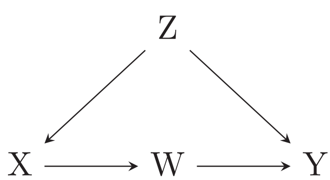

Figure 4. Causal diagram appropriate for front-door method application.

Chapter 7 gives multiple examples. One such example looks at a case of sequential decision-making, where one (randomized) treatment leads to an endogneous intermediate outcome, which itself determines whether a second treatment is applied. Identifying the dual effect of both treatments is straightforward in the *do*-calculus, but is difficult with standard methods, and this case seems likely to have multiple applications in economics.

3.4. Equilibrium models

The previous three subsections detailed areas where economists are likely to find a lot of marginal benefit from learning Pearl's approach to causal inference. The next two subsections examine topics common in economics where, even if SCM is capable of handling these topics (which it always is; recall that any theorem in PO can be converted to one in SCM), it seems unlikely that economists would find it worthwhile to go out of their way to apply SCM.

The first of these is in the case of equilibrium models like supply and demand, or other models with simultaneous equations or bi-directional causality. These models occur regularly in economics, and economists are familiar with a number of ways to identify parameters in these models, including instrumental variables.

These sorts of causal inference problems cannot be solved by causal diagrams alone (except sometimes in cases where instrumental variables are sufficient – see *The Book of Why* Chapter 7), and someone learning SCM from *The Book of Why* alone would not be able to address them. Standard causal diagrams must be *acyclic*, whereas equilibrium models contain feedback loops and so are cyclic. The difficulty that standard causal diagrams have with equilibrium models was a particular concern in the Imbens (2020) review of SCM.

It is possible to construct causal diagrams with cycles and simultaneity, but identification in these diagrams is a more advanced subject relative to the way diagrams are introduced in *The Book of Why* (a short discussion of, and introduction to, identification in cyclic diagrams is in Section 18.6 of Pearl, 2022). Identification in these diagrams requires a return to the structural equations and the solution of nonlinear equations, much like economists are already used to in solving these systems.

3.5. Identification by functional form

Causal diagrams are, by their design, nonparametric. Because of this, causal diagrams are unable to fully demonstrate any approach to identification that relies on assumptions about functional form, for example regression discontinuity designs, or any design that relies on effect moderation or an interaction term in its regression specification. SCM as a whole can incorporate functional form restrictions, but it must be done in the structural equations. The functional form assumption necessary for identification is 'hidden' on a different layer of the SCM and cannot be seen on the diagram.

This is a problem, as such designs are common in applied econometrics, as are conditions like monotonicity in instrumental variables designs. Attempts to use causal diagrams to demonstrate these methods sometimes do so by breaking the rules of causal diagrams and including terms representing functional form Cunningham (2021) and Huntington-Klein (2021).

It is possible that causal diagrams can be used to aid the process of covariate selection in these models (although this may lead to error in cases where the functional form assumption itself renders a covariate unnecessary). However, the diagram obscures the actual research design for a reader, and most of the work with functional form assumptions in SCM must be done in the structural equations, which may not add much beyond what economists already do in these cases. It seems unlikely that economists will find marginal value in SCM in application to these designs.

4. Conclusion

The Book of Why provides an introductory view to Judea Pearl's contributions to causal inference. Becoming acquainted with Pearl's approach may be valuable if only to converse with the fields making use of it. But beyond that, SCM offers some interesting advantages to economists performing causal inference, at least in some areas.

Economists are already familiar with the concept of structural modeling, and while Pearl's use of structural equations is similarly useful, economists already have access to structural modeling tools, or are at least aware of them and have chosen not to use them. The marginal benefit of Pearl's approach to structural equations in areas where structural modeling is already in place, or where causal diagrams are not themselves much of a marginal contribution, is relatively low, although difficult structural identification problems may still get additional benefit from the *do*-calculus. As Heckman and Pinto (2022) point out, however, in cases where econometric structural modeling is already used, causal diagrams and the *do*-calculus are a simplified representation of the structural model that lose some powers of identification relative to that structural model, and may not be able to address economic content as commonly-applied as the Roy model. Pearl would point out that the same problem could be addressed with a full set of structural equations in SEM, but again the marginal benefit relative to econometric structural modeling is not as apparent in this case.

For that reason, the greatest marginal benefit to economists of learning Pearl's system are the applications for which causal diagrams on their own are the most useful, and in cases where economists who do not currently use explicit structural modeling may be able to impose structure via the diagram itself, such as in many reduced-form or quasiexperimental applications. Valuable contributions include the processes of covariate selection and the development of placebo tests, both of which are highly important parts of causal inference that are currently done in economics on a somewhat ad-hoc basis.

The Book of Why, which as an introductory text does not extend to the full use of structural equation modeling, is an excellent introduction to the philosophical and practical benefits of using Pearl's SCM approach. Economists are likely to find value in thinking through problems using causal diagrams, and at least in application to some problems, they are likely to find that bringing diagrams into their toolbox will be worthwhile.

Notes

1. Those familiar with the history of causal inference in economics are likely to bristle at reading *The Book of Why*, where one gets the clear sense, although never explicitly stated and perhaps unintended, that the book views causal inference before Pearl's SCM approach as, at best, a half-born preview of Pearl's work. We can look past this.
2. This distinction can be summed up by looking at the Pearl (2020) response to Imbens (2020). Even in the cases where this paper and Imbens come to a very similar conclusion, I do not think Pearl's response to Imbens applies in the same way to this paper, given my applied-researcher target and marginal-benefit question. This does not mean that Pearl would endorse my findings; given (Pearl, 2020) he would likely find issue with what my paper implies can be achieved without SCM, although among economists this would be uncontroversial.
3. Note that the third step here is presented in the simplified form it appears as in *The Book of Why* and becomes more complex in the presence of other interventions or conditioning variables. Heiss (2021) offers a presentation of the *do*-calculus that is between *The Book of Why* and Pearls' more technical materials that provides a fuller picture of operation 3 and should be accessible to economists. Example derivations for 'front-door adjustment' (discussed below) can be found in *The Book of Why* Chapter 7, and for 'back-door adjustment' (selection on observables) in Heiss (2021).

Disclosure statement

No potential conflict of interest was reported by the author.

Notes on contributor

Nick Huntington-Klein is an assistant professor of economics at Seattle University. He is known for his work in higher education policy and metascience, as well as for his accessible materials on causal inference.

References

Bellemare, M. F., Bloem, J. R., & Wexler, N. (2020). The paper of how: Estimating treatment effects using the front-door criterion [Technical report].

Cinelli, C., Forney, A., & Pearl, J. (2020). A crash course in good and bad controls. SSRN 3689437.

Cunningham, S. (2021). *Causal inference: The mixtape*. Yale University Press.

Currie, J., Kleven, H., & Zwiers, E. (2020). Technology and big data are changing economics: Mining text to track methods. In *AEA papers and proceedings* (Vol. 110, pp. 42–48). American Economic Association.

Galles, D., & Pearl, J. (1998). An axiomatic characterization of causal counterfactuals. *Foundations of Science*, *3*(1), 151–182. https://doi.org/10.1023/A:1009602825894

Heckman, J. J., & Pinto, R. (2022). Causality and econometrics [Technical report]. National Bureau of Economic Research.

Heiss, A. (2021, September). Do-calculus adventures! https://www.andrewheiss.com/blog/2021/09/07/do-calculus-backdoors/

Huntington-Klein, N. (2021). *The effect: An introduction to research design and causality*. Chapman and Hall/CRC.

Imbens, G. W. (2020). Potential outcome and directed acyclic graph approaches to causality: Relevance for empirical practice in economics. *Journal of Economic Literature*, *58*(4), 1129–1179. https://doi.org/10.1257/jel.20191597

Morgan, S. L., & Winship, C. (2015). *Counterfactuals and causal inference*. Cambridge University Press.

Pearl, J. (2020, January). On imbens's comparison of two approaches to empirical economics. https://causality.cs.ucla.edu/blog/index.php/category/do-calculus/

Pearl, J. (2022). Causal diagrams for empirical research (with discussions). In *Probabilistic and causal inference: The works of Judea Pearl* (pp. 255–316). Association for Computing Machinery.

Pearl, J., & Mackenzie, D. (2018). *The book of why: The new science of cause and effect*. Basic Books.

Rubin, D. B. (1974). Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of Educational Psychology*, *66*(5), 688. https://doi.org/10.1037/h0037350

Spirtes, P., Glymour, C. N., Scheines, R., & Heckerman, D. (2000). *Causation, prediction, and search*. MIT Press.

Textor, J., Hardt, J., & Knüppel, S. (2011). DAGitty: A graphical tool for analyzing causal diagrams. *Epidemiology*, *22*(5), 745. https://doi.org/10.1097/EDE.0b013e318225c2be

White, H., & Lu, X. (2011). Causal diagrams for treatment effect estimation with application to efficient covariate selection. *Review of Economics and Statistics*, *93*(4), 1453–1459. https://doi.org/10.1162/REST_a_00153

Wright, S. (1921). Correlation and causation. *Journal of Agricultural Research*, *20*, 557–585.