Title

Causal models and learning from data: integrating causal modeling and statistical estimation.

Journal

Epidemiology (Cambridge, Mass.), 25(3)

ISSN

1044-3983

Authors

Petersen, Maya L
van der Laan, Mark J

Publication Date

2014-05-01

DOI

10.1097/ede.0000000000000078

Peer reviewed

Published in final edited form as:
Epidemiology. 2014 May ; 25(3): 418–426. doi:10.1097/EDE.0000000000000078.

Causal Models and Learning from Data:

Integrating Causal Modeling and Statistical Estimation

Maya L. Petersen and **Mark J. van der Laan**
Divisions of Biostatistics and Epidemiology, University of California, Berkeley, School of Public Health, Berkeley, CA

Abstract

The practice of epidemiology requires asking causal questions. Formal frameworks for causal inference developed over the past decades have the potential to improve the rigor of this process. However, the appropriate role for formal causal thinking in applied epidemiology remains a matter of debate. We argue that a formal causal framework can help in designing a statistical analysis that comes as close as possible to answering the motivating causal question, while making clear what assumptions are required to endow the resulting estimates with a causal interpretation. A systematic approach for the integration of causal modeling with statistical estimation is presented. We highlight some common points of confusion that occur when causal modeling techniques are applied in practice and provide a broad overview on the types of questions that a causal framework can help to address. Our aims are to argue for the utility of formal causal thinking, to clarify what causal models can and cannot do, and to provide an accessible introduction to the flexible and powerful tools provided by causal models.

Epidemiologists must ask causal questions. Describing patterns of disease and exposure is not sufficient to improve health. Instead, we seek to understand why such patterns exist and how we can best intervene to change them. The crucial role of causal thinking in this process has long been acknowledged in our field's historical focus on confounding.

Major advances in formal causal frameworks have occurred over the past decades. Several specific applications, such as the use of causal graphs to choose adjustment variables[1] or the use of counterfactuals to define the effects of longitudinal treatments,[2] are now common in the epidemiologic literature. However, the tools of formal causal inference have the potential to benefit epidemiology much more extensively.

We argue that the wider application of formal causal tools can help frame sharper scientific questions, make transparent the assumptions required to answer these questions, facilitate rigorous evaluation of the plausibility of these assumptions, clearly distinguish the process of causal inference from the process of statistical estimation, and inform analyses of data and interpretation of results that rigorously respect the limits of knowledge. We, together

Correspondence: Maya L. Petersen, University of California, Berkeley, 101 Haviland Hall, Berkeley, CA 94720-7358. mayaliv@berkeley.edu.

The authors report no conflicts of interest.

with others, advocate for a systematic approach to causal questions that involves (1) specification of a causal model that accurately represents knowledge and its limits; (2) specification of the observed data and their link to the causal model; (3) translation of the scientific question into a counterfactual quantity; (4) assessment of whether, under what assumptions, this quantity is identified–whether it can be expressed as a parameter of the observed data distribution or estimand; (5) statement of the resulting statistical estimation problem; (6) estimation, including assessment of statistical uncertainty; and (7) interpretation of results (Figure 1).[3,4] We emphasize how causal models can help navigate the ubiquitous tension between the causal questions posed by public health and the inevitably imperfect nature of available data and knowledge.

A GENERAL ROADMAP FOR CAUSAL INFERENCE

1. Specify knowledge about the system to be studied using a causal model

Of the several models available, we focus on the structural causal model,[5–10] which provides a unification of the languages of counterfactuals,[11,12] structural equations,[13,14] and causal graphs.[1,7] Structural causal models provide a rigorous language for expressing both background knowledge and its limits.

Causal graphs represent one familiar means of expressing knowledge about a data-generating process; we focus here on directed acyclic graphs (Figure 2). Figure 2A provides an example of a directed acyclic graph for a simple data-generating system consisting of baseline covariates *W*, an exposure *A*, and an outcome *Y*. Such graphs encode causal knowledge in several ways. First, graphs encode knowledge about the possible causal relations among variables. Knowledge that a given variable is not directly affected by a variable preceding it is encoded by omitting the corresponding arrow (referred to as “an exclusion restriction”). Figure 2A reflects the absence of any such knowledge; baseline covariates *W* may have affected the exposure *A*, and both may have affected the outcome *Y*. In other cases, investigators may have knowledge that justifies exclusion restrictions. For example, if *A* represents adherence to a randomly assigned treatment *R* (Figure 2B), it might be reasonable to assume that random assignment (if effectively blinded) had no effect on the outcome other than via adherence. Such knowledge is represented by omission of an arrow from *R* to *Y*. Second, omission of a double-headed arrow between two variables assumes that any unmeasured “background factors” that go into determining the values that these variables take are independent (or, equivalently, that the variables do not share an unmeasured cause, referred to as “an independence assumption”). Figure 2A makes no independence assumptions, whereas Figure 2B reflects the knowledge that, because *R* was randomly assigned, it shares no unmeasured common cause with any other variables.

The knowledge encoded in a causal graph can alternatively be represented using a set of structural equations, in which each node in the graph is represented as a deterministic function of its parents and a set of unmeasured background factors. The error term for a given variable *X* (typically denoted as U_X) represents the set of unmeasured background factors that, together with variable *X*’s parents (nodes with arrows pointing to *X*), determine what value *X* takes (Figure 2). The set of structural equations, together with any restrictions placed on the joint distribution of the error terms (expressed on the graph as assumptions

about the absence of unmeasured common causes between two nodes) together constitute a structural causal model.[5,6]

Such a structural causal model provides a flexible tool for encoding a great deal of uncertainty about the true data-generating process. Specifically, a structural causal model allows for uncertainty about the existence of a causal relationship between any two variables (through inclusion of an arrow between them), uncertainty about the distribution of all unmeasured background factors that go into determining the value of these variables (frequently, no restrictions are placed on the joint distribution of the errors, beyond any independence assumptions), and uncertainty about the functional form of causal relationships between variables (the structural equations can be specified nonparametrically). If knowledge in any of these domains is available, however, it can be readily incorporated. For example, if it is known that R was assigned independently to each subject with probability 0.5, this knowledge can be reflected in the corresponding structural equation (Figure 2B), as can parametric knowledge on the true functional form of causal relationships.

In sum, the flexibility of a structural causal model allows us to avoid many (although not all) unsubstantiated assumptions and thus facilitates specification of a causal model that describes the true data-generating process. Alternative causal models differ in their assumptions about the nature of causality and make fewer untestable assumptions.[15–19]

2. Specify the observed data and their link to the causal model

Specification of how the observed data were generated by the system described in our causal model provides a bridge between causal modeling and statistical estimation.

The causal model (representing knowledge about the system to be studied) must be explicitly linked to the data measured on that system. For example, a study may have measured baseline covariates W, exposure A, and outcome Y on an independent random sample of n individuals from some target population. The observed data on a given person thus consist of a single copy of the random variable $O = (W, A, Y)$. If our causal model accurately describes the data-generating system, the data can be viewed as n independent and identically distributed draws of O from the corresponding system of equations.

More complex links between the observed data and the causal model are also possible. For example, study participants may have been sampled on the basis of exposure or outcome status. More complex sampling schemes such as these can be handled either by specifying alternative links between the causal model and the observed data or by incorporating selection or sampling directly into the causal model.[6,20]

The structural causal model is assumed to describe the system that generated the observed data. This assumption may or may not have testable implications. For example, the systems described by Figures 2A, D, and E could generate any possible distribution of $O = (W, A, Y)$. We thus say that these causal models place no restrictions on the joint distribution of the observed data, implying a nonparametric statistical model. In contrast, the system described by Figure 2B can generate only distributions of $O = (W, R, A, Y)$ in which R is independent

of W (a testable assumption). This causal model thus implies a semiparametric statistical model. Independence (and conditional independence) restrictions of this nature can be read from the graph using the criterion of d-separation.[5,7] (The set of possible distributions for the observed data may also be restricted via functional or inequality constraints.)[21–23]

The statistical model should reflect true knowledge about the data-generating process, ensuring that it contains the true distribution of the observed data. Although in some cases parametric knowledge about the data-generating process may be available, in many cases a causal model that accurately represents knowledge is compatible with any possible distribution for our observed data, implying a nonparametric statistical model.

3. Specify the target causal quantity

The formal language of counterfactuals forces explicit statement of a hypothetical experiment to answer the scientific question of interest.

Specification of an ideal experiment and a corresponding target counterfactual quantity helps ensure that the scientific question drives the design of a data analysis and not vice versa.[24] This process forces the researcher to define exactly which variables would ideally be intervened on, what the interventions of interest would look like, and how the resulting counterfactual outcome distributions under these interventions would be compared.

A causal model on the counterfactual distributions of interest can be specified directly (and need not be graphical)[2,11,12,25] or, alternatively, it can be derived by representing the counterfactual intervention of interest as an intervention on the graph (or set of equations). The initial structural causal model describes the set of processes that could have generated (and thus the set of possible distributions for) the observed data. The postintervention causal model describes the set of processes that could have generated the counterfactual variables we would have measured in our ideal experiment and thus the set of possible distributions for these variables. Figure 2C provides an illustration of the postintervention structural causal model corresponding to Figure 2A, under an ideal experiment in which exposure A is set to 0 for all persons.

One common counterfactual quantity of interest is the average treatment effect: the difference in mean outcome that would have been observed had all members of a population received versus not received some treatment. This quantity is expressed in terms of counterfactuals as $E(Y_1-Y_0)$, where Y_a denotes the counterfactual outcome under an intervention to set $A = a$. The corresponding ideal experiment would force all members of a population to receive the treatment and then roll back the clock and force all to receive the control.

In many cases, the ideal experiment is quite different. For example, if the exposure of interest were physical exercise, one might be interested in the counterfactual risk of mortality if all subjects without a contraindication to exercise were assigned the intervention (an example of a “realistic” dynamic regime, in which the treatment assignment for a given person depends on that person’s characteristics).[26–28] Furthermore, if the goal is to evaluate the impact of a policy to encourage more exercise, a target quantity might compare the

existing outcome distribution with the distribution of the counterfactual outcome under an intervention to shift the exercise distribution, while allowing each person's exercise level to remain random (an example of a stochastic intervention).[17,29–31] Additional examples include effects of interventions on multiple nodes and mediation effects.[2,32–36] Figure 3 lists major decision points when specifying a counterfactual target parameter and provides examples of general categories of causal questions that can be formally defined using counterfactuals.

4. Assess identifiability

A structural causal model provides a tool for understanding whether background knowledge, combined with the observed data, is sufficient to allow a causal question to be translated into a statistical estimand, and, if not, what additional data or assumptions are needed.

Step 3 translated the scientific question into a parameter of the (unobserved) counterfactual distribution of the data under some ideal intervention(s). We say that this target causal quantity is identified, given a causal model and its link to the observed data, if the target quantity can be expressed as a parameter of the distribution of the observed data alone—an estimand. Structural causal models provide a general tool for assessing identifiability and deriving estimands that, under explicit assumptions, equal causal quantities.[5,37–39]

A familiar example is provided by the use of causal graphs to choose an adjustment set when estimating the average treatment effect (or other parameter of the distribution of Y_a). For observed data consisting of n independent identically distributed copies of $O = (W, A, Y)$, when preintervention covariates W block all unblocked back-door paths from A to Y in the causal graph (the "back-door criterion"),[5,7] then the distribution of Y_a is identified according to the "G-computation formula" (given in Equation 1 for discrete valued O).[6] The same result also holds under the "randomization assumption" that Y_a is independent of A given W[18]:

$$P(Y_a=y)=\sum_w P(Y=y|A=a,\ W=w)P(W=w) \quad (1)$$

Equation 1 equates a counterfactual quantity (the left-hand side) with an estimand (the right-hand side) that can be targeted for statistical estimation. The back-door criterion can be straightforwardly evaluated using the causal graph; it fails in Figure 2A and B but holds in Figure 2D and E. Single world intervention graphs allow graphical evaluation of the randomization assumption for a given causal model.[17]

Although the estimand in Equation 1 might seem to be an obvious choice, use of a formal causal framework can be instrumental in choosing an estimand. For example, in Figure 2F,

$$P\left(Y_a=y\right) \neq \sum_{wz} P(Y=y|A=a,\ W=w,\ Z=z)P(W=w,\ Z=z).$$

The use of a structural causal model in this case warns against adjustment for *Z*, even if it precedes *A*. In other cases, adjustment for interventions that occur after the exposure may be warranted.

Many scientific questions imply more complex counter-factual quantities (Figure 3), for which no single adjustment set will be sufficient and alternative identifiability results are needed. Causal frameworks provide a tool for deriving these results, often resulting in new estimands and thereby suggesting different statistical analyses. Examples include effect mediation,[19,32–34] the effects of interventions at multiple time points,[2,18,41] dynamic interventions,[17,18,42] causal and non-causal parameters in the presence of informative censoring or selection bias,[2,6,20] and the transport of causal effects to new settings.[27,43–45]

5. Commit to a statistical model and estimand

A causal model that accurately represents knowledge can help to select an estimand as close as possible to the wished-for causal quantity, while emphasizing the challenge of using observational data to make causal inferences.

In many cases, rigorous application of a formal causal framework forces us to conclude that existing knowledge and data are insufficient to claim identifiability—in itself a useful contribution. The process can also often inform future studies by suggesting ways to supplement data collection.[46] However, in many cases, better data are unlikely to become available or “current best” answers are needed to questions requiring immediate action.

One way to navigate this tension is to rigorously differentiate assumptions that represent real knowledge from assumptions that do not, but which if true, would result in identifiability. We refer to the former as “knowledge-based assumptions” and the latter as “convenience-based assumptions.” An estimation problem that aims to provide a current best answer can then be defined by specifying: (1) a statistical model implied by knowledge-based assumptions alone (and thus ensured to contain the truth); (2) an estimand that is equivalent to the target causal quantity under a minimum of convenience-based assumptions; and (3) a clear differentiation between convenience-based assumptions and real knowledge. In other words, a formal causal framework can provide a tool for defining a statistical estimation problem that comes as close as possible to addressing the motivating scientific question, given the data and knowledge currently available, while remaining transparent regarding the additional assumptions required to endow the resulting estimate with a causal interpretation.

For example, measured variables are rarely known to be sufficient to control confounding. Figure 2A may represent the knowledge-based causal model, under which the effect of *A* on *Y* is unidentified; however, under the augmented models in Figure 2D and E, result (1) would hold. If our goal is to estimate the average treatment effect, we might thus define a statistical estimation problem in which: (1) the statistical model is nonparametric (as implied by Figure 2A and by Figure 2D and E); and (2) we select

$$\sum_{w}(E(Y|A=1,\ W=w)-E(Y|A=0,\ W=w))P(W=w) \quad (2)$$

as the estimand. We can then proceed with estimation of Equation 2, while making explicit the conditions under which the estimand may diverge from the causal effect of interest.

6. Estimate

Choice between estimators should be motivated by their statistical properties.

Once the statistical model and estimand have been defined, there is nothing causal about the resulting estimation problem. A given estimand, such as Equation 2, can be estimated in many ways. Estimators of Equation 2 include those based on inverse probability weighting,[47] propensity score matching,[40] regression of the outcome on exposure and confounders (followed by averaging with respect to the empirical distribution of confounders), and double robust efficient methods,[48,49] including targeted maximum likelihood.[4]

To take another example, marginal structural models (Figure 3) are often used to define target causal quantities, particularly when the exposure has multiple levels.[2,35] Under causal assumptions, such as the randomization assumption (or its sequential counterpart), this quantity is equivalent to a specific estimand. However, once the estimand has been specified, estimation itself is a purely statistical problem. The analyst is free to choose among several estimators; inverse probability–weighted estimators are simply one popular class.

Any given class of estimator itself requires, as "ingredients," estimators of specific components of the observed data distribution. For example, one approach to estimating Equation 2 is to specify an estimator of $E(Y|A, W)$. ($P(W = w)$ is typically estimated as the sample proportion.) In many cases, the true functional form of this conditional expectation is unknown. In some cases, $E(Y|A, W)$ can be estimated non-parametrically using a saturated regression model; however, at moderate sample sizes or if W or A are high dimensional or contain continuous covariates, the corresponding contingency tables will have empty cells and this approach will break down. As a result, in many practical scenarios, the analyst must either rely on a parametric regression model that is likely to be misspecified (risking significant bias and misleading inference) or do some form of dimension reduction and smoothing to trade off bias and variance.

Similar considerations arise when specifying an inverse probability–weighted or propensity score–based estimator. In this case, estimator consistency depends on consistent estimation of the treatment mechanism or propensity score $P(A = a|W)$. An extensive literature on data-adaptive estimation addresses how best to approach this tradeoff for the purposes of fitting a regression object such as $E(Y|A, W)$ or $P(A = a|W)$.[50] An additional literature on targeted estimation discusses how to update resulting estimates to achieve an optimal bias-variance tradeoff for the final estimand, as well as how to generate valid estimates of statistical uncertainty when using data-adaptive approaches.[4]

In sum, there is nothing more or less causal about alternative estimators. However, estimators have important differences in their statistical properties, and these differences can result in meaningful differences in performance in commonly encountered scenarios, such as

strong confounding.[51–53] Choice among estimators should be driven by their statistical properties and by investigation of performance under conditions similar to those presented by a given applied problem.

7. Interpret

Use of a causal framework makes explicit and intelligible the assumptions needed to move from a statistical to a causal interpretation of results.

Imposing a clear delineation between knowledge-based assumptions and convenience-based assumptions provides a hierarchy of interpretation for an analysis (Figure 4). The use of a causal framework to define a problem in no way undermines purely statistical interpretation of results. For example, an estimate of Equation 2 can be interpreted as an estimate of the difference in mean outcome between exposed and unexposed subjects who have the same values of observed baseline covariates (averaged with respect to the distribution of the covariates in the population). The same principle holds for more complex analyses—for example, if inverse probability weighting is used to estimate the parameters of a marginal structural model.

The use of a formal causal framework ensures that the assumptions needed to augment the statistical interpretation with a causal interpretation are explicit. For example, if we believe that Figure 2D or E represents the true causal structure that generated our data, then our estimate of Equation 2 can be interpreted as an estimate of the average treatment effect. The use of a causal model and a clear distinction between convenience-based and knowledge-based assumptions makes clear that the true value of an estimand may differ from the true value of the causal effect of interest. The magnitude of this difference is a causal quantity; choice of statistical estimator does not affect it. Statistical bias in an estimator can be minimized through data-driven methods, while evaluating the likely or potential magnitude of the difference between the estimand and the wished-for causal quantity requires alternative approaches (sometimes referred to as sensitivity analyses).[54–58]

There is nothing in the structural causal model framework that requires the intervention to correspond to a feasible experiment.[6,59,60] In particular, some parameters of substantive interest (such as the natural direct effect) cannot be defined in terms of a testable experiment but can nonetheless be defined and identified under a structural causal model.[32] If, in addition to the causal assumptions needed for identifiability, the investigator is willing to assume that the intervention used to define the counterfactuals corresponds to a conceivable and well-defined intervention in the real world, interpretation can be further expanded to include an estimate of the impact that would be observed if that intervention were to be implemented in practice (in an identical manner, in an identical population, or, under additional assumptions, in a new population).[27,43–45] Finally, if the intervention could be feasibly randomized with complete compliance and follow-up, interpretation of the estimate can be further expanded to encompass the results that would be seen if one implemented this hypothetical trial. The decision of how far to move along this hierarchy can be made by the investigator and reader based on the specific application at hand. The assumptions required are explicit and, when expressed using a causal graph, readily understandable by subject matter experts.

The debate continues as to whether causal questions and assumptions should be restricted to quantities that can be tested and thereby refuted via theoretical experiment.[15–17,19,61] Given that most such experiments will be done, if ever, only after the public health relevance of an analysis has receded,[17,19] the extent to which such a restriction will enhance the practical impact of applied epidemiology seems unclear. However, while we have chosen to focus on the structural casual model, our goal is not to argue for the supremacy of a single class of causal model as optimal for all scientific endeavor. Rather, we suggest that the systematic application of a roadmap will improve the impact of applied analyses, irrespective of formal causal model chosen.

CONCLUSIONS

Epidemiologists continue to debate whether and how to integrate formal causal thinking into applied research. Many remain concerned that the use of formal causal tools leads both to overconfidence in our ability to estimate causal effects from observational data and to the eclipsing of common sense by complex notation and statistics. As Levine articulates this position, “the language of ‘causal modeling’ is being used to bestow the solidity of the complex process of causal inference upon mere statistical analysis of observational data… conflating statistical and causal inference.”[62] We argue that a formal causal framework, when used appropriately, represents a powerful tool for preventing exactly such conflation and overreaching interpretation.[6,63]

Like any tool, the benefits of a causal inference framework depend on how it is used. The ability to define complex counterfactual parameters does not ensure that they address interesting or relevant questions. The ability to formally represent causal knowledge as graphs or equations is not a license to exaggerate what we know. The ability to prove equivalence between a target counterfactual quantity and an estimand under specific assumptions does not make these assumptions true, nor does it ensure that they can be readily evaluated. The best estimation tools can still produce unreliable statistical estimates when data are inadequate.

Good epidemiologic practice requires us to learn as much as possible about how our data are generated, to be clear about our question, to design an analysis that answers this question as well as possible using available data, to avoid or minimize assumptions not supported by knowledge, and to be transparent and skeptical when interpreting results. Used appropriately, a formal causal framework provides an invaluable tool for integrating these basic principles into applied epidemiology. We argue that the routine use of formal causal modeling would improve the quality of epidemiologic research, as well as research in the innumerable disciplines that aim to use statistics to learn how the world works.

Acknowledgments

M.L.P. is a recipient of a Doris Duke Clinical Scientist Development Award. M.J.v.d.L. is supported by NIH award R01 AI074345.

References

1. Greenland S, Pearl J, Robins JM. Causal diagrams for epidemiologic research. Epidemiology. 1999; 10:37–48. [PubMed: 9888278]
2. Robins JM, Hernán MA, Brumback B. Marginal structural models and causal inference in epidemiology. Epidemiology. 2000; 11:550–560. [PubMed: 10955408]
3. Pearl, J. The Causal Foundations of Structural Equation Modeling. In: Hoyle, RH., editor. Handbook of Structural Equation Modeling. New York: Guilford Press; 2012. p. 68-91.
4. van der Laan, M.; Rose, S. Targeted Learning: Causal Inference for Observational and Experimental Data. Berlin, Heidelberg, New York: Springer; 2011.
5. Pearl, J. Causality: Models, Reasoning, and Inference. New York: Cambridge University Press; 2000.
6. Pearl J. Causal diagrams for empirical research. Biometrika. 1995; 82:669–710.
7. Pearl J. Causal inference in statistics: an overview. Statistics Surveys. 2009; 3:96–146.
8. Strotz RH, Wold HO. Recursive vs. nonrecursive systems: an attempt at synthesis (part I of a triptych on causal chain systems). Econometrica. 1960; 28:417–427.
9. Haavelmo T. The statistical implications of a system of simultaneous equations. Econometrica. 1943; 11:1–12.
10. Wright S. Correlation and causation. J Agric Res. 1921; 20:557–585.
11. Neyman J. On the application of probability theory to agricultural experiments. Essay on principles. Section 9. Stat Sci. 1923; 5:465–472.
12. Rubin D. Estimating causal effects of treatments in randomized and non-randomized studies. J Educational Psychol. 1974; 66:688–701.
13. Duncan, O. Introduction to Structural Equation Models. New York: Academic Press; 1975.
14. Goldberger A. Structural equation models in the social sciences. Econometrica: Journal of the Econometric Society. 1972; 40:979–1001.
15. Spirtes, P.; Glymour, C.; Scheines, R. Causation, Prediction, and Search. Number 81 in Lecture Notes in Statistics. New York/Berlin: Springer-Verlag; 1993.
16. Dawid AP. Causal inference without counterfactuals. J Am Stat Association. 2000; 95:407–424.
17. Richardson, TS.; Robins, JM. Technical Report 128. Seattle, WA: Center for Statistics and the Social Sciences, University of Washington; 2013. Single World Intervention Graphs (SWIGs): A Unification of the Counterfactual and Graphical Approaches to Causality.
18. Robins J. A new approach to causal inference in mortality studies with a sustained exposure period. Mathem Mod. 1986; 7:1393–512.
19. Robins, JM.; Richardson, T. Alternative graphical causal models and the identification of direct effects. In: Shrout, P.; Keyes, K.; Ornstein, K., editors. Causality and psychopathology: Finding the determinants of disorders and their cures. Oxford, UK: Oxford University Press; 2010. p. 103-58.
20. Hernán MA, Hernández-Díaz S, Robins JM. A structural approach to selection bias. Epidemiology. 2004; 15:615–625. [PubMed: 15308962]
21. Kang, C.; Tian, J. Inequality Constraints in Causal Models with Hidden Variables. Arlington, VA: AUAI Press; 2006.
22. Shpitser, I.; Pearl, J. Proceedings of the 23rd National Conference on Artificial Intelligence. Vol. 2. Menlo Park, CA: AAAI Press; 2008. Dormant independence; p. 1081-1087.
23. Shpitser, I.; Richardson, TS.; Robins, JM.; Evans, R. [Accessed October 28, 2013.] Parameter and Structure Learning in Nested Markov Models. UAI Workshop on Causal Structure Learning. 2012. Available at: http://www.stat.washington.edu/tsr/uai-causal-structure-learning-workshop/
24. Hernan MA, Alonso A, Logan R, et al. Observational studies analyzed like randomized experiments: an application to postmenopausal hormone therapy and coronary heart disease. Epidemiology. 2008; 19:766–779. [PubMed: 18854702]
25. Robins, JM. Structural nested failure time models. In: Armitage, P.; Colton, T., editors. Encyclopedia of Biostatistics. Chichester, UK: John Wiley and Sons; 1998. p. 4372-4389.

26. Hernán MA, Lanoy E, Costagliola D, Robins JM. Comparison of dynamic treatment regimes via inverse probability weighting. Basic Clin Pharmacol Toxicol. 2006; 98:237–242. [PubMed: 16611197]
27. Robins J, Orellana L, Rotnitzky A. Estimation and extrapolation of optimal treatment and testing strategies. Stat Med. 2008; 27:4678–4721. [PubMed: 18646286]
28. van der Laan MJ, Petersen ML. Causal effect models for realistic individualized treatment and intention to treat rules. Int J Biostat. 2007; 3:Article 3. [PubMed: 19122793]
29. Taubman SL, Robins JM, Mittleman MA, Hernán MA. Intervening on risk factors for coronary heart disease: an application of the parametric g-formula. Int J Epidemiol. 2009; 38:1599–1611. [PubMed: 19389875]
30. Muñoz ID, van der Laan M. Population intervention causal effects based on stochastic interventions. Biometrics. 2012; 68:541–549. [PubMed: 21977966]
31. Cain LE, Robins JM, Lanoy E, Logan R, Costagliola D, Hernán MA. When to start treatment? A systematic approach to the comparison of dynamic regimes using observational data. Int J Biostat. 2010; 6:Article 18. [PubMed: 21972433]
32. Pearl, J. Direct and indirect effects. Proceeedings of the 17th Conference on Uncertainty in Artificial Intelligence; San Francisco, CA: Morgan Kaufmann; 2001. p. 411-420.
33. Petersen ML, Sinisi SE, van der Laan MJ. Estimation of direct causal effects. Epidemiology. 2006; 17:276–284. [PubMed: 16617276]
34. Robins JM, Greenland S. Identifiability and exchangeability for direct and indirect effects. Epidemiology. 1992; 3:143–155. [PubMed: 1576220]
35. Robins J. Association causation and marginal structural models. Synthese. 1999; 121:151–179.
36. Robins, J.; Hernan, MA. Estimation of the causal effects of time-varying exposures. In: Fitzmaurice, G.; Davidian, M.; Verbeke, G.; Molenbergh, G., editors. Longitudinal Data Analysis. London: Chapman and Hall/CRC; 2009. p. 566-568.
37. Shpitser I, Pearl J. Complete identification methods for the causal hierarchy. J Mach Learn Res. 2008; 9:1941–1979.
38. Tian, J.; Pearl, J. Proceedings of the 18th National Conference on Artificial Intelligence. Menlo Park, CA: AAAI Press; 2002. A general identification condition for causal effects; p. 567-573.
39. Tian, J.; Shpitser, I. On identifying causal effects. In: Dechter, R.; Geffner, H.; Halpern, J., editors. Heuristics, Probability and Causality: A Tribute to Judea Pearl. UK: College Publications; 2010. p. 415-444.
40. Rosenbaum P, Rubin D. The central role of the propensity score in observational studies. Biometrika. 1983; 70:41–55.
41. Robins J. Addendum to: “A new approach to causal inference in mortality studies with a sustained exposure period—application to control of the healthy worker survivor effect”. Comput Math Appl. 1987; 14:923–945.
42. Tian, J. Identifying dynamic sequential plans. Proceeedings of the 24th Conference on Uncertainty in Artificial Intelligence; Helsinki, Finland: AUAI Press; 2008. p. 554-561.
43. Petersen ML. Compound treatments, transportability, and the structural causal model: the power and simplicity of causal graphs. Epidemiology. 2011; 22:378–381. [PubMed: 21464653]
44. Hernán MA, VanderWeele TJ. Compound treatments and transportability of causal inference. Epidemiology. 2011; 22:368–377. [PubMed: 21399502]
45. Pearl, J.; Bareinboim, E. Transportability Across Studies: A Formal Approach. Los Angeles, CA: Computer Science Department, University of California; 2011.
46. Geng EH, Glidden DV, Bangsberg DR, et al. A causal framework for understanding the effect of losses to follow-up on epidemiologic analyses in clinic-based cohorts: the case of HIV-infected patients on antiretroviral therapy in Africa. Am J Epidemiol. 2012; 175:1080–1087. [PubMed: 22306557]
47. Hernán MA, Robins JM. Estimating causal effects from epidemiological data. J Epidemiol Community Health. 2006; 60:578–586. [PubMed: 16790829]
48. Robins J, Rotnitzky A, Zhao L. Estimation of regression coefficients when some regressors are not always observed. J Am Stat Assoc. 1994; 89:846–866.

49. Scharfstein DO, Rotnitzky A, Robins JM. Adjusting for nonignorable drop-out using semiparametric nonresponse models. J Am Stat Assoc. 1999; 94:1096–120.
50. Hastie, T.; Tibshirani, R.; Friedman, J. The Elements of Statistical Learning: Data Mining, Inference, and Prediction. 2. New York: Springer-Verlag; 2009.
51. Kang J, Schafer J. Demystifying double robustness: a comparison of alternative strategies for estimating a population mean from incomplete data. Stat Science. 2007; 22:523–539.
52. Petersen ML, Porter KE, Gruber S, Wang Y, van der Laan MJ. Diagnosing and responding to violations in the positivity assumption. Stat Methods Med Res. 2012; 21:31–54. [PubMed: 21030422]
53. Moore KL, Neugebauer R, Laan MJ, Tager IB. Causal inference in epidemiological studies with strong confounding. Statistics in Medicine. 2012; 31:1380–404. [PubMed: 22362629]
54. Robins, J.; Rotnitzky, A.; Scharfstein, D. Sensitivity analysis for selection bias and unmeasured confounding in missing data and causal inference models. In: Halloran, M.; Berry, D., editors. Statistical Models in Epidemiology: The Environment and Clinical Trials. New York, NY: Springer; 1999. p. 1-92.
55. Imai K, Keele L, Yamamoto T. Identification, inference, and sensitivity analysis for causal mediation effects. Stat Sci. 2010; 25:51–71.
56. Greenland S. Multiple-bias modelling for analysis of observational data. J Royal Stat Soc A. 2005; 168:267–306.
57. Vanderweele TJ, Arah OA. Bias formulas for sensitivity analysis of unmeasured confounding for general outcomes, treatments, and confounders. Epidemiology. 2011; 22:42–52. [PubMed: 21052008]
58. Díaz I, van der Laan MJ. Sensitivity analysis for causal inference under unmeasured confounding and measurement error problems. Int J Biostat. 2013; 9:149–160. [PubMed: 24246288]
59. Hernán MA. Invited commentary: hypothetical interventions to define causal effects–afterthought or prerequisite? Am J Epidemiol. 2005; 162:618–620. discussion 621. [PubMed: 16120710]
60. van der Laan M, Haight T, Tager I. Respond to " Hypothetical interventions to define causal effects". Am J Epidemiol. 2005; 162:621–622.
61. Pearl, J. [Accessed October 28, 2013.] Causal analysis in theory and practice: on mediation, counterfactuals and manipulations. Posting to UCLA Causality Blog. May 3. 2010 Available at: http://www.mii.ucla.edu/causality/?p=133. UCLA Causality Blog2010
62. Levine B. Causal models [letter]. Epidemiology. 2009; 20:931. [PubMed: 19829191]
63. Hogan JW. Causal Models: the author responds. Epidemiology. 2009; 20:931–932. [PubMed: 19829191]
64. Petersen ML, Wang Y, van der Laan MJ, Guzman D, Riley E, Bangsberg DR. Pillbox organizers are associated with improved adherence to HIV antiretroviral therapy and viral suppression: a marginal structural model analysis. Clin Infect Dis. 2007; 45:908–915. [PubMed: 17806060]
65. Bodnar LM, Davidian M, Siega-Riz AM, Tsiatis AA. Marginal structural models for analyzing causal effects of time-dependent treatments: an application in perinatal epidemiology. Am J Epidemiol. 2004; 159:926–934. [PubMed: 15128604]
66. Hernán MA, McAdams M, McGrath N, Lanoy E, Costagliola D. Observation plans in longitudinal studies with time-varying treatments. Stat Methods Med Res. 2009; 18:27–52. [PubMed: 19036915]
67. Rosenblum M, Jewell NP, van der Laan M, Shiboski S, van der Straten A, Padian N. Analyzing direct effects in randomized trials with secondary interventions: an application to HIV prevention trials. J R Stat Soc Ser A Stat Soc. 2009; 172:443–465.
68. Gsponer T, Petersen M, Egger M, et al. The causal effect of switching to second-line ART in programmes without access to routine viral load monitoring. AIDS. 2012; 26:57–65. [PubMed: 22089376]
69. Petersen ML, van der Laan MJ, Napravnik S, Eron JJ, Moore RD, Deeks SG. Long-term consequences of the delay between virologic failure of highly active antiretroviral therapy and regimen modification. AIDS. 2008; 22:2097–2106. [PubMed: 18832873]
70. Westreich D, Cole SR, Tien PC, et al. Time scale and adjusted survival curves for marginal structural Cox models. Am J Epidemiol. 2010; 171:691–700. [PubMed: 20139124]

71. Wester CW, Stitelman OM, deGruttola V, Bussmann H, Marlink RG, van der Laan MJ. Effect modification by sex and baseline CD4+ cell count among adults receiving combination antiretroviral therapy in Botswana: results from a clinical trial. AIDS Res Hum Retroviruses. 2012; 28:981–988. [PubMed: 22309114]

72. Shpitser, I.; Pearl, J. Effects of treatment on the treated: identification and generalization. Proceeedings of the 24th Conference on Uncertainty in Artificial Intelligence; Arlington, VA: AUAI Press; 2009. p. 514-521.

1. Specify knowledge about the system to be studied using a causal model.

Represent background knowledge about the system to be studied. A causal model describes the set of possible data-generating processes for this system.

2. Specify the observed data and their link to the causal model.

Specify what variables have been or will be measured, and how these variables are generated by the system described by the causal model.

3. Specify a target causal quantity.

Translate the scientific question into a formal causal quantity (defined as some parameter of the distribution of counterfactual random variables), following the process in Figure 3.

4. Assess identifiability.

Assess whether it is possible to represent the target causal quantity as a parameter of the observed data distribution (estimand), and, if not, what further assumptions would allow one to do so.

5. State the statistical estimation problem.

Specify the estimand and statistical model. If knowledge is sufficient to identify the causal effect of interest, commit to the corresponding estimand. If not, but one still wishes to proceed, choose an estimand that under minimal additional assumptions would equal or approximate the causal effect of interest.

6. Estimate.

Estimate the target parameter of the observed data distribution, respecting the statistical model.

7. Interpret.

Select among a hierarchy of interpretations, ranging from purely statistical to approximating a hypothetical randomized trial.

FIGURE 1.

A general roadmap for approaching causal questions.

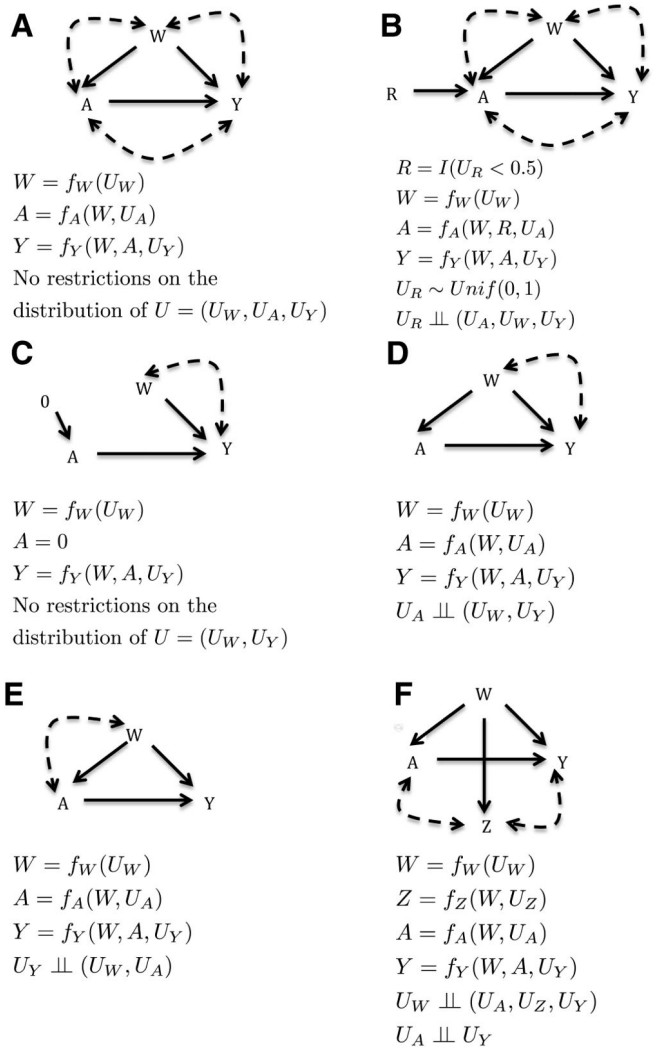

FIGURE 2.

Causal graphs and corresponding structural equations representing alternative data generating processes. (A), Assumes that W may have affected A, both A and W may have affected Y, and W, A, and Y may all share unmeasured common causes. (B), Assumes that subjects were randomly assigned to an exposure arm R with probability 0.5 and that assigned exposure arm R does not affect the outcome Y other than via uptake of the exposure A. (C), Figure 2A after a hypothetical intervention in which all subjects in the target population are assigned exposure level 0 (however this is defined). (D) and (E), Figure 2A augmented with additional assumptions to ensure identifiability of the distribution of Y_a (and thus target quantities such as the average treatment effect) given observed data (W,A,Y). (F), Assumes that Z does not affect A or Y, that A shares no common causes with W or Y, and that W shares no common causes with Y or Z. Figures 2A, 2D, and 2E have no testable implications with W,A,Y observed, and thus imply non-parametric statistical models for the distribution of (W,A,Y). In contrast, Figure 2B implies that R is independent of W, and thus implies a semi-parametric statistical model for the distribution of (W,R,A,Y).

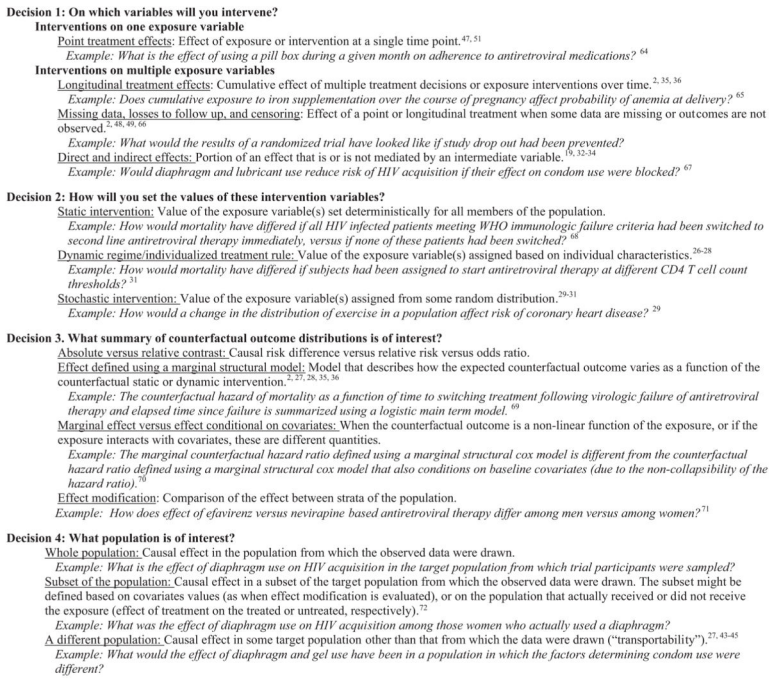

FIGURE 3.

Translating a scientific question into a target causal quantity.

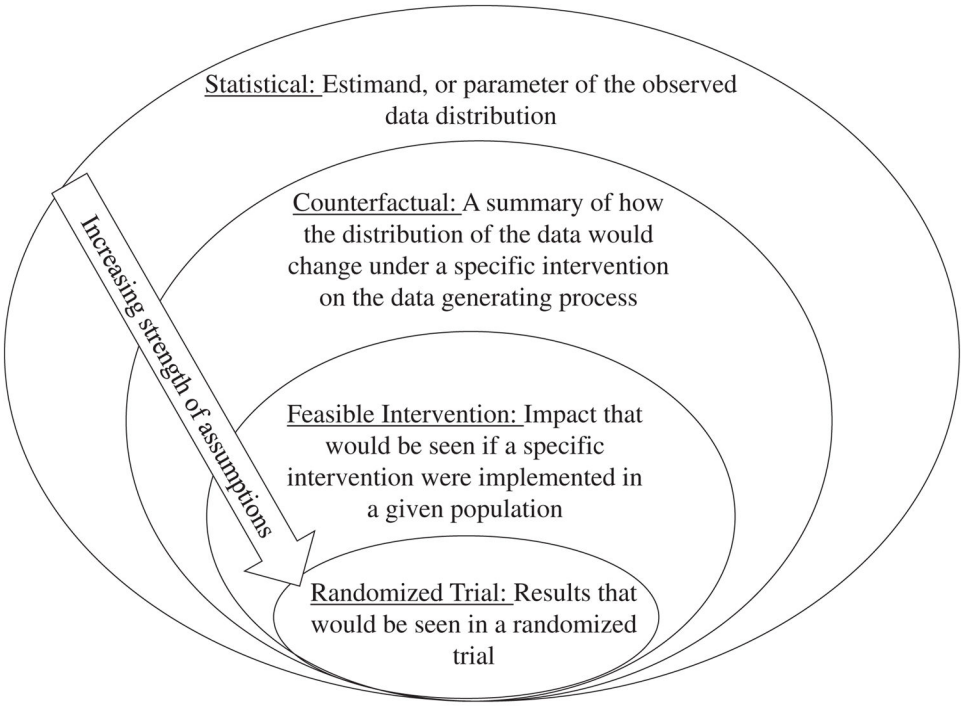

FIGURE 4.

Interpreting the results of a data analysis: a hierarchy. The use of a statistical model known to contain the true distribution of the observed data and of an estimator that minimizes bias and provides a valid measure of statistical uncertainty helps to ensure that analyses maintain a valid statistical interpretation. Under additional assumptions, this interpretation can be augmented.