Children and Youth Services Review 35 (2013) 1130–1142

Contents lists available at ScienceDirect

Children and Youth Services Review

journal homepage: www.elsevier.com/locate/childyouth

Child welfare and the challenge of causal inference

E. Michael Foster [a,*], Kimberly McCombs-Thornton [b]

[a] Department of Health Care Organization and Policy, School of Public Health, University of Alabama at Birmingham, Birmingham, AL, United States
[b] Department of Maternal and Child, Gillings School of Global Public Health, University of North Carolina, Chapel Hill, NC, United States

ARTICLE INFO

Article history:
Received 11 March 2011
Accepted 17 March 2011
Available online 26 March 2011

Keywords:
Causal inference
Propensity scores
Methodology

ABSTRACT

Causal inference refers to the assessment of cause and effect relationships in observational data—i.e., in situations where random assignment is impossible or impractical. Choices involving children in the child welfare system evoke core elements of causal inference—manipulation and the counterfactual. How would a child's circumstances differ if we changed her environment? This article begins with one mathematical approach to framing causal inference, the potential outcomes framework. This methodology is a cornerstone of newer approaches to causal inference. This framework makes clear the identification problem inherent in causal inference and highlights a key assumption often used to identify the model (ignorability or no unobserved confounding). The article then outlines the various approaches to causal inference and organizes them around whether they assume ignorability as well as other key features of each approach. The article then provides guidelines for producing good causal inference. These guidelines emerge from a review of methodological literature as broad ranging as epidemiology, statistics, economics, and policy analysis. These steps will be illustrated using an example from child welfare. The article concludes with suggestions for how the field could apply these newer methods.

1. Introduction

Causal inference refers to the assessment of cause and effect relationships in observational data—i.e., in situations where random assignment is impossible or impractical. One might want to know, for example, the consequences of a risk factor, such as maternal substance use, on a child's development. Obviously one cannot randomly place children with mothers who are using alcohol or other drugs. Similar questions arise as to key choices by child welfare system staff such as whether to remove a child from her home. Presumably that choice is made based on the child's best interest in the eyes of the case worker and others; to simply assign a child to another situation would be unethical. As a result, observational data represent the main source of information that can inform child welfare policy, and for that reason, causal inference lies at the heart of research on children and youth services.

Choices involving children in the child welfare system evoke core elements of causal inference—manipulation and the counterfactual (Pearl, 2000). How would a child's circumstances differ if we changed her environment? The fundamental challenge of causal inference is that at a point in time, we observe children only in one condition or circumstance.[1] Causal inference is the question of "what if"—what if the child's circumstances were different? In many cases, the question is what if social policy changed those circumstances?

More formally, the quantity we want to calculate is the difference between what we observe and the counterfactual we do not. We have two unknowns (what is and what might have been) and one known (what we observe); statisticians refer to such a situation as "under-identified". How does one identify a model? Generally, identification requires either an assumption of some sort (e.g., the absence of confounding by an unobserved characteristic) and/or additional data (Heckman & Vytlacil, 2007b). The latter often involves a comparison group of some sort—the former specifies the conditions under which the experiences of that group represent the counterfactual we need. Generally, some combination of the two is required. For example, we might observe the treated entity prior to treatment, and that information is valuable. However, causal inference requires an assumption as well—some assumption about other forces acting on the entity over time.

In the case of maternal substance abuse, a natural comparison group would involve the children of other women, and one easily could compute differences in outcomes between the two groups. A lay person would recognize that such a comparison is potentially misleading—children in the two groups differ in many ways, and addressing the needs of the substance-abusing mother would only partially eliminate those differences. The association is likely much larger than the true effect, and for that reason, every graduate student (hopefully) is taught that "association does not mean causation" (Pearl, 2009).

* Corresponding author. Tel.: +1 919 259 6466.
E-mail address: Mike4kids@gmail.com (E.M. Foster).

[1] Economists refer to this challenge as the "evaluation problem" (Heckman & Vytlacil, 2007a). As this reference also notes, not all agree that random assignment is the gold standard nor that a potential cause need be manipulable.

0190-7409/$ – see front matter © 2011 Elsevier Ltd. All rights reserved.
doi:10.1016/j.childyouth.2011.03.012

The number of poor examples of causal inference, however, indicates that this lesson is either not absorbed or not appreciated fully.[2] Indeed, *moving from association to causation is audacious*: it is a venture into the unknown.[3] Most statistical tools have mechanical properties that refine associations in one way or another. Regression, for example, is the orthogonal projection of an outcome onto the space spanned by the covariates. In non-technical terms, regression imposes the restriction that none of the covariates are correlated with the error term in the model *whether that is actually true or not.* The fundamental problem of causal inference dictates that for the resulting estimate to be causal, other assumptions need to be added—namely, that treatment is not simultaneously determined with the outcome nor is it correlated with unobserved determinants of the outcome.[4] Many of the key assumptions are untestable. (A reasonable rule of thumb is that the more essential an assumption is to causal inference, the less likely it is that one can test it. Regression, for example, often assumes a linear relationship between the covariates and the outcome. This assumption can be tested but is certainly not required for causal inference.)[5]

Causal inference may be hard, but the good news is that methodologists are developing new tools. This development is refining—and narrowing—the assumptions under which one can move from association to causation. It is also clarifying which assumptions embedded in our analytical techniques are required for causal inference and which serve some other purposes. In many instances, the latter can be relaxed through better empirical practice. For example, outliers may damage causal inference, but many regression diagnostics are available to identify such problems. The bottom line, though, is that a key assumption (ignorability, described below) will remain even if the application of regression is improved.

An informal discussion highlights key issues, but a full treatment requires a mathematical presentation. Only such a presentation can provide the specificity needed. For example, such specificity makes it clear that some causal questions require assumptions that others do not. For example, the assumption required to assess the effect of teenage childbearing on teen mothers differs from that required to assess the effect on other young women *were they to make that choice.*

This article begins with one mathematical approach to frame causal inference, the potential outcomes framework (Holland, 1986; Rubin, 2005). This methodology is a cornerstone of newer approaches to causal inference. This framework makes clear the identification problem inherent in causal inference and highlights a key assumption often used to identify the model (ignorability or no unobserved confounding). The article then outlines the various approaches to causal inference and organizes them around whether they make this assumption as well as other key features of each approach. The article then provides guidelines for producing causal inference. These guidelines emerge from a review of methodological literature as broad ranging as epidemiology, statistics, economics, and policy analysis. These steps will be illustrated using an example from child welfare. It concludes with suggestions for how the field could apply these newer methods.

A brief word about terminology is in order before proceeding. If not already apparent, the term "treatment" is used very generally in this literature. It could refer to any actual treatment or service but also to exposures or conditions or states. There may be very little about these conditions that are beneficial.

2. Potential outcomes framework

2.1. Basic mathematical language

This basic mathematical language provides specificity about what we really want to know and what we have to assume to estimate that quantity. In the potential outcomes framework, one is interested in a treatment (or exposure or characteristic) (D) and some outcome (Y). D = 1 or 0 for the treated and untreated groups, respectively. For each individual, one can think of her as having two possible outcomes Y_1 (outcome when treated, D = 1) and Y_0 (outcome when not treated, D = 0). One can characterize each individual by three variables, (Y_1, Y_0, and D)—the outcome if treated, if not treated and treatment status. The fundamental problem is that we do not observe (the joint distribution of) all three random variables. Rather we observe D for each individual (was she treated or not) and either Y_1 or Y_0 for the treated (D = 1) and untreated (D = 0), respectively. One can write $Y_{obs} = (D^*Y_1) + (1-D)^*Y_0$, where Y_{obs} is the observed outcome. Economists refer to this relationship as the "the observation rule".

We have omitted a person subscript to this point, but these outcomes differ across individuals, both because individuals differ in the treatment they receive and in other factors (differentiating even those who receive the same treatment). What we would like to know is the treatment effect—i.e., $\tau = Y_{1,i} - Y_{0,i}$—for every individual i. Generally, we are unable to calculate this term for a single person.

Can we get traction on this problem if we reduce our goal? What if we wanted to know only $\tau = E[Y_{1,i} - Y_{0,i}]$, the mean of this distribution of effects? The "E" identifies the expectations operator or the average. When one writes E[Y|X = x] or E[Y|x], this notation refers to the average value of Y for a specific value or range (x) of X. (Following convention, random variables are denoted with upper case variables and specific values of those variables, lower case.)

Because we chose the mean and not some other characteristic of the distribution (such as the median), one can see that $\tau = E[Y_{1,i}] - E[Y_{0,i}]$.[6] The problem is (still) that neither of the two terms is observed for the whole population. In particular, we observe not $E[Y_{1,i}]$ but $E[Y_{1,i}|D_i = 1]$ and not $E[Y_{0,i}]$ but $E[Y_{0,i}|D_i = 0]$. One sees the outcome under treatment only for those treated and the converse for untreated individuals. These quantities are related to the terms of interest, the expectation of the outcome in the treated state for all persons in the sample:

$$E\left[Y_{1,i}\right] = p(D_i = 1)E\left[Y_{1,i}|D_i = 1\right] + p(D_i = 0)\mathbf{E}\left[\mathbf{Y_{1,i}}|\mathbf{D_i} = \mathbf{0}\right] \qquad (1)$$

The average outcome for all individuals is the weighted average of the treated outcome for the treated and untreated. The weights are the proportions of the sample that do and do not receive treatment. The emboldened term is unobserved—that expectation and the expression as a whole cannot be calculated without some additional assumption. Again, this is the fundamental problem of causal inference.

[2] The poor examples of causal inference are too many to count. For an example, compare Christakis, Zimmerman, DiGiuseppe, & McCarty (2004) with Foster & Watkins (2010) and see discussion in Foster (2010a).

[3] We recognize the difficulty of this assessment in our lives—what if I had gone to graduate school in public health instead of economics and so on—and we recognize how difficult that question is to answer. Why this question would be easier to answer about anonymous people about whom we have less information is difficult to fathom. But in our personal lives we often recognize that the question is unanswerable: "Who knows?" That hesitation also should inform efforts to gage causal relationships in observational data.

[4] Heckman and Vytlacil (2007b) make this distinction in a different way. Regression conditions on a variable mechanically; whether it "fixes" it in a causal way—i.e., reveals its influence if manipulated—requires additional assumptions. This distinction is also apparent in Pearl's well known treatment (Pearl, 2000). He represents the latter with the do() operator. This operator represents the effect of making the exposure determined by forces outside of the model.

[5] This discussion oversimplifies matters somewhat. The importance of different assumptions also depends on the nature of the causal question. The literature in statistics focuses on the non-parametric estimation of the ATT, ATE and ATU, and for that purpose, linearity (or a specific functional form) is not required. As Heckman and Vytlacil (2007a) notes, this is only one of several causal questions related to policy issues. These questions include estimating the impact of a program involving combinations of treatments that have not yet been delivered.

[6] An added advantage of working with the expectation is that this estimand does not require knowledge of the correlation between Y_0 and Y_1. That correlation is required for calculation of other potential treatment effects of interest, such as the proportion of individuals who benefit from treatment.