Children and Youth Services Review 35 (2013) 1143–1151

Contents lists available at ScienceDirect

Children and Youth Services Review

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

Causal effects of foster care: An instrumental-variables approach ☆

Joseph J. Doyle Jr. *

MIT Sloan School of Management, 77 Massachusetts Avenue, E62-515, Cambridge, MA 02139, United States

ARTICLE INFO

Article history:
Received 25 August 2010
Accepted 11 March 2011
Available online 26 March 2011

Keywords:
Foster care
Instrumental variables
Causal effects

ABSTRACT

This paper describes the use of instrumental-variables (IV) to estimate causal effects of foster care on long- and short-term outcomes. This estimation strategy provides a tool to evaluate what are known as "natural experiments": settings that mimic randomization usually associated with a controlled trial. The proposed natural experiment involves the effective randomization of investigators to child-protection cases. The results suggest that foster care placement increases like likelihood of delinquency and emergency healthcare episodes. Care must be taken when interpreting IV estimates. The results apply to cases that are part of the natural experiment—"marginal cases" where the investigators may disagree about the placement recommendation.

1. Introduction

There is no dispute that severely abused or neglected children should be protected, and a foster family home has been judged the best alternative whenever possible. A key policy question is one of degree: how aggressive should child protective services be? Child protection agencies trade off two competing goods: family preservation and child protection (Barth, 1999; Lindsey, 1994; Maluccio, Pine, & Warsh, 1994). More aggressive child protection may reduce child abuse or neglect, but removal from parents may be traumatic to children as well. For example, much has been written about the potential for such instability to hinder child development, and multiple placements once a child has been placed in foster care has been associated with greater emotional and behavioral problems among foster children.[1]

A better understanding of the causal effects of foster care on short- and long-term outcomes for children at risk of placement would be useful to inform child-welfare policy. These effects are difficult to estimate because of confounding factors (Testa & Poertner, 2010; Vinnerljung, Sundell, Andree Lofholm & Eva, 2006; Courtney, 2000; Gelles, 2000; Goerge, Wulczyn & Fanshel, 1994; Jonson-Reid & Barth, 2000; National Research Council, 1998; McDonald, Allen, Westerfelt & Piliavin, 1996).[2] The main estimation problem is that children placed in foster care likely differ from children who remain at home. In particular, worse outcomes for foster children compared to other children in the same area could be due to abusive family backgrounds, as opposed to any effect of foster care placement (Kerman, Wildfire, & Barth, 2002). Indeed, foster care policy directly targets children who appear to be at high risk of poor life outcomes. Former foster children are far more likely than are others to drop out of school, be imprisoned, enter the homeless population, join welfare, or experience substance abuse problems (Clausen, Landsverk, Ganger, Chadwick & Litrownik, 1998; Courtney & Piliavin, 1998; Dworsky & Courtney, 2000; US DHHS, 1999).

To estimate causal effects of a given treatment on outcomes, it would be useful to conduct a randomized, controlled trial. Such a trial is unrealistic when the treatment is foster care placement. Another approach uses naturally-occurring randomization to mimic that of a trial. The current paper builds on earlier work (Doyle, 2007, 2008), where the source of randomization comes from the rotational assignment of cases to child-protection investigators.[3] One family

☆ Special thanks to Joshua Angrist Mark Duggan, Steve Levitt, Robert Goerge, Michael Greenstone, Jon Gruber, Jim Poterba, Tom Stoker, Mark Testa, Roberto Rigobon, Tavneet Suri, and Robert Moffitt for comments and advice on this research program. I would like to acknowledge the Chapin Hall Center for Children at the University of Chicago for the creation of the Integrated Database on Child and Family Programs in Illinois that was used in this study. All findings, interpretations and conclusions based on the use of the IDB are solely my responsibility and do not necessarily represent the views of the Chapin Hall Center for Children. I would also like to acknowledge the generous support of the National Science Foundation under grant SES-0518757.

* Tel.: +1 617 452 3761.
E-mail address: jjdoyle@mit.edu.

[1] There is a large empirical literature on placement stability, as it is one observable characteristic in administrative data. See James, Landsverk and Slymen (2004), Newton, Litronwnik and Landsverk (2000), and Smith, Stormshak, Chamberlain and Whaley (2001).

[2] Few studies compare children investigated for abuse. See Runyan and Gould (1985), Elmer (1986), Davidson-Arad, Englechin-Segal and Wozner (2003), and Wald, Carlsmith and Leiderman (1988) for four small scale studies. Jonson-Reid and Barth (2000) studied 160,000 children in California using administrative data and found lower delinquency on average for children who remained at home, especially those who received in-home services.

[3] The main goal of this paper is to describe the instrumental-variable strategy. More details about the natural experiment and other nuances can be found in the earlier papers. Further, similar instrumental-variable strategies have been used in other settings. See, for example, Kling (2006) who studies the random assignment of judges to estimate effects of prison-sentence length on labor-market outcomes.

0190-7409/$ – see front matter
doi:10.1016/j.childyouth.2011.03.014

may be assigned an investigator that is more likely to recommend placement, and the next family reported to that field office may be assigned a different investigator who is less likely to do so.

The goal of this paper is to demonstrate how instrumental-variable techniques can be used to measure causal effects in a natural-experimental setting. Two outcomes are considered: juvenile delinquency later in life and emergency healthcare usage within the year following the abuse report. Meanwhile, the empirical strategy relies on two main estimates: the extent to which the investigator assigned to the case is associated with (1) foster care placement and (2) the outcomes of interest—juvenile delinquency and emergency healthcare. If the only way that the investigators affect children is through foster care placement, then the instrumental-variable strategy combines these estimates to investigate causal effects of foster care. Part of the paper discusses how these results should be interpreted. Namely, the results apply to children on the margin of placement—those cases where investigators may disagree about the placement recommendation.

The paper is organized as follows. Section 2 presents the empirical framework. Section 3 describes the background in support of the natural experiment as a useful strategy. Section 4 describes the data, and section five presents the instrumental-variable examples. Section 6 concludes.

2. Empirical framework

This section briefly describes the use of instrumental variables to estimate causal effects.[4] The example considered here is the effect of foster care on long- and short-term outcomes. For the sake of this section, the outcome of interest will be juvenile delinquency: does placement in foster care increase or reduce the likelihood that a child will enter the juvenile justice system and by how much? For example, a parameter of interest is given by:

$$(\text{Juvenile Deliquency} \mid \text{Foster Care} = 1) \\ -\text{E}(\text{Juvenile Delinquency} \mid \text{Foster Care} = 0) \tag{1a}$$

The discussion first considers a simple mean comparison across individuals: some were placed in foster care while others were not. It then incorporates controls for observable differences in the two groups in a regression framework and in a propensity-score matching framework. Last, instrumental-variables estimation is discussed in the context of these estimators.

2.1. Naïve estimate: mean comparison

One estimate uses the sample analogs of (1): a difference in means for individuals who were in foster care compared to others who were never placed. Take a group of adults and estimate the following regression for individual i:

$$\text{JD}_i = \beta_0 + \beta_1 \text{FC}_i + \varepsilon_i. \tag{1b}$$

This equation is referred to as the *structural equation*. JD is an indicator taking a value of one if individual i was a juvenile delinquent later in adolescence and zero otherwise; FC is similar but reflects whether the individual was placed in foster care. The estimate of β_1 is the mean difference in delinquency across the two groups.

In this simple model, every individual has the same relationship between foster care placement and juvenile delinquency given by β_1. This assumption of a common coefficient could be relaxed to allow the coefficient to vary across individuals, becoming β_{1i}. The parameter of interest (1) would describe the average of these β_1's, known as the "average treatment effect".

The main concern with this type of comparison is that foster care is not randomly assigned, and it seems likely that there are omitted variables in the statistical model. The foster care indicator is likely related to ε: the factors that are related to juvenile delinquency but not in the model. For example, family background characteristics such as child abuse or neglect can affect the likelihood of foster care placement and the underlying propensity of juvenile delinquency (Widom, 1989). A higher delinquency rate among former foster children may reflect the underlying abuse or neglect rather than an effect of foster care *per se*.

2.2. Conditional expectation

A related approach to estimate the effect of foster care on delinquency would add control variables, X. The parameters of interest then could depend on X:

$$\text{E}(\text{Juvenile Deliquency} \mid \text{Foster Care} = 1, X) \\ -\text{E}(\text{Juvenile Deliquency} \mid \text{Foster Care} = 0, \text{X}). \tag{2a}$$

One way to incorporate these controls is to add them to the estimating equation:

$$\text{JD}_i = \beta_0 + \beta_1 \text{FC}_i + \beta_2 \text{X}_i + \varepsilon_i \tag{2b}$$

where X is a vector of control variables. Ordinarily least squares (OLS) provides an estimate of 2 using variation in foster care placement that unrelated—or "orthogonal"—to the characteristics in X, such as measures of child abuse allegations. The idea is to consider variation in FC that is unrelated to variation stemming from the vector of observable characteristics, X, when estimating β_1.

A concern with such a conditional expectation is that it is not possible to observe in datasets the same characteristics observed by those who decided on the foster care placement. Investigators and judges use practice wisdom to arrive at a conclusion based on factors that are difficult to quantify to include in a statistical model (Cash, 2001). This suggests that statistical models will omit key variables, as they are not available in the data.

2.3. Flexible controls: propensity score and other matching estimators

The functional form in Eq. (2b) can be relaxed, for example by estimating Eq. (1b) separately for particular case characteristics. This is known as a matching estimator, and the estimates from these cells could be of independent interest or aggregated to calculate an average causal response.

When there are many controls to consider, few individuals may have with the same set of covariates.[5] A popular form of matching aggregates these covariates into a "propensity score" (Rosenbaum & Rubin, 1983; Ryan, Hong, Herz, and Hernandez, 2010). First, the likelihood that a child is placed in foster care would be estimated with all of the control variables. This provides a predicted propensity for every individual based on her particular observable characteristics. Second, individuals with similar propensities are compared.

One approach is to estimate Eq. (2a) for different subsets of the propensity score, such as deciles. This can be a useful way to describe the data and investigate heterogeneity of effects across different children. Typically, the observable characteristics will be shown to be similar, or "balanced", across individuals who received the treatment and those that did not within these deciles. The assumption is that the unobserved characteristics are similar as well, and, thus, can be

[4] More formal reviews include Angrist & Krueger, 2001; Heckman, Lalonde, & Smith, 1999. Textbook treatments include Angrist & Pischke, 2009; Wooldridge, 2002; Cameron & Trivedi, 2005, and for a similar exposition of the use of instrumental variables but with a healthcare example, see McClellan, McNeil & Newhouse, 1994.

[5] There is a tradeoff between a closely-matched comparison and the limited sample size leading to imprecise estimates.