Causal Effects of Foster Care: An Instrumental-Variables Approach

Joseph J. Doyle, Jr.*
MIT Sloan School of Management & NBER

January 2011

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.

Key Words: Foster care; Instrumental Variables; Causal Effects

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, and 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 and Poertner, forthcoming; Vinnerljung et al., 2006, Courtney, 2000; Gelles, 2000; Goerge, Wulzcyn, and Fanshel, 1994; Jonson-Reid and Barth, 2000; National Research Council, 1998; McDonald et al., 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, and Barth, 2002). Indeed, foster care policy

---

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

[2] Few studies compare children investigated for abuse. See Runyan and Gould (1985), Elmer (1986), Davidson-Arad, et al. (2003), and Wald et al. (1998) 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.

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 et al., 1998; Courtney and Piliavin, 1998; Dworsky and 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; Doyle 2008), where the source of randomization comes from the rotational assignment of cases to child-protection investigators.[3] One family 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

---

[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.

margin of placement—those cases where investigators may disagree about the placement recommendation.

The paper is organized as follows. Section two presents the empirical framework. Section three describes the background in support of the natural experiment as a useful strategy. Section four describes the data, and section five presents the instrumental-variable examples. Section six 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:

$(1a)$ E(Juvenile Deliquency | Foster Care = 1) - E(Juvenile Delinquency | Foster Care = 0)

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.A. Naïve estimate: Mean comparison

---

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

One estimate uses the sample analogues 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:

$$(1b) \ \text{JD}_i = \beta_0 + \beta_1 \text{FC}_i + \varepsilon_i$$

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.B. 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:

$(2a)$ E(Juvenile Deliquency | Foster Care =1, X) - E(Juvenile Deliquency | Foster Care =0, X)

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

$$(2b) \text{ JD}_i = \beta_0 + \beta_1 \text{FC}_i + \beta_2 X_i + \varepsilon_i$$

where X is a vector of control variables. Ordinarily least squares (OLS) provides an estimate of β_1 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.C. Flexible controls: Propensity score and other matching estimators

The functional form in (2b) can be relaxed, for example by estimating (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 and Rubin, 1983, Ryan et al., forthcoming). 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 (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 ignored.[6] This is an assumption, however, and like the concern about omitted variables in equation (2b), differences in unobserved characteristics within these matched individuals can continue to confound the estimates.

2.D. Instrumental variables

As noted in the introduction, an instrumental variable, Z, describes a natural experiment. There are two main assumptions that make an instrumental variable strategy possible.

Assumption 1 (First Stage): Z is related to FC

---

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

[6] This is known as the ignorability assumption. More formally, the potential outcomes if an individual were or were not placed in foster care are assumed to be independent of the assignment into foster care, conditional on the covariates. For more detail, see Rosenbaum and Rubin (1983) and Heckman, Ichimura, and Todd (1998).

The first ingredient is that Z is associated with the treatment variable of interest. The relationship is known as the first stage and is testable by estimating the relationship between the instrument and the treatment:

$$(3)\ \mathrm{FC}_i = \alpha_0 + \alpha_1 Z_i + \mu_i$$

α_1 can be estimated by OLS, and a strong level of statistical significance and plausible sign and magnitude are sought. This can also be estimated with more than one instrument, in which case Z is a vector. This model can also be augmented to include control variables as well.

Assumption 2 (Exclusion Restriction): Z is not related to ε

The second assumption is sometimes called the exclusion restriction, meaning that Z can be excluded from the structural equation (1b).[7] This means that the variation in FC due to Z is unrelated to the unobserved characteristics that can affect the outcome of interest (juvenile delinquency).

Just as the relationship between FC and ε in (1b) cannot be tested, neither can the relationship between Z and ε. If Z were randomly assigned, then Assumption 2 would follow. Often random assignment is not possible, however, and institutional detail—such as changes in policy regimes or administrative assignment procedures—may be used to describe natural experiments where Assumption 2 is plausible.

One piece of evidence that can serve as a check on the plausibility of Assumption 2 is to see whether the observable characteristics, X, are related to Z. If they were, it would be less convincing that the unobservable characteristics happen to be unrelated to Z. This is similar to

---

[7] Z is uncorrelated with ε: $\mathrm{Cov}(Z, \varepsilon)=0$ (Wooldridge, 2002).

checking to see if the observable characteristics are balanced for matching estimators. In the end, concerns that unobserved characteristics are related to Z remain. A judgment has to be made as to whether the natural experiment is successful in leading to treatment that is "as good as random" or not.

To consider how the instrument works, we can substitute the first stage (3) into the structural equation (1b). This reduces the system from two to one equation, which is referred to as the "reduced form":

$$(4)\ \mathrm{JD}_i = \beta_0 + \beta_1(\alpha_0 + \alpha_1 Z_i + \mu_i) + \varepsilon_i$$

$$= (\beta_0 + \beta_1\alpha_0) + \beta_1\alpha_1 Z_i + +(\beta_1\mu_i + \varepsilon_i)$$

$$= \gamma_0 + \gamma_1 Z_i + \vartheta_i$$

Note that if Z is unrelated to ε, then an OLS estimation of (4) will yield an unbiased estimate of γ_1. At times it is easier to defend the reduced form as a parameter of interest. For example, if the instrument relates to policy variation, then γ_1 is an estimate of the effect of that policy change, regardless of whether the effect only goes through the treatment of interest, in this case FC.

2.E. Mechanics 1: Wald Estimator

When the instrument is binary—Z takes on a value of 0 or 1—the instrumental-variable estimator can be written as a Wald estimator:

$$(5)\ \frac{\mathrm{E}(\mathrm{JD}{=}1|Z{=}1) - \mathrm{E}(\mathrm{JD}{=}1|Z{=}0)}{\mathrm{P}(\mathrm{FC}{=}1|Z{=}1) - \mathrm{E}(\mathrm{FC}{=}1|Z{=}0)}$$

The numerator is the reduced form estimate of γ_1, which is $\beta_1\alpha_1$. The denominator is the first stage relationship between the instrument and the treatment of interest, α_1. By dividing the two, one arrives at an estimate of β_1.[8]

The Wald representation demonstrates how the instrumental-variable technique works. The variation in juvenile delinquency is coming from the instrument, shown in the numerator as the instrument taking on a value of one compared to zero. Likewise, the variation in foster care is coming from the instrument.

If the instrument affects the child only through its effect on foster care placement, then the change in the numerator would come from this channel. Further, if the change in the outcome variable described by the numerator is caused by a change in the propensity of foster care that is less than one, then the denominator essentially re-scales the reduced form to arrive at the causal effect if foster care indicator changed from zero to one. For example, if assignment to a strict investigator increased the probability of going into foster care by 0.5, then the change in delinquency would have to stem from this change in the proportion placed in care. To estimate the effect of foster care placement itself—a change in the probability of foster care from zero to one—then the reduced-form estimate would need to be doubled. This is what the Wald estimator does, by dividing by 0.5. The example below focuses on a binary instrument to relate back to this formulation. For illustration, this type of calculation will be shown in the empirical results.

2.F. Mechanics 2: Two-Stage Least Squares (2SLS)

---

[8] Technically, the estimation of the first stage results in IV estimates that are biased in small samples but converge to β_1 in large samples.

The instrumental-variable strategy is often estimated by two-stage least squares (2SLS). The second stage estimates the relationship between the predicted value of FC from the first stage, (3), given by $\widetilde{FC}$:

$$(6) \ \ \text{JD}_i = \beta_0 + \beta_1 \widetilde{FC}_i + \varepsilon_i$$

This can be estimated by OLS, although in practice, the first and second stages (equations 3 and 6) can be estimated simultaneously in most statistical software packages—a calculation that corrects the standard errors for the fact that $\widetilde{FC}$ is estimated. The above mechanics of plugging the first stage estimate into the second stage equation is again meant to demonstrate that the estimates are coming from the variation in foster care associated with the instrument.

As in the mean comparisons, the instrumental variable strategy can be carried out with control variables. Here, the first and second stages would both be estimated with the same control variables.[9]

If Assumption 2 does not hold, then the estimate will not converge to β_1 even in large samples. The bias is given by the strength of the relationship between Z and ε (Assumption 2) divided by the strength of the relationship between Z and X (Assumption 1).[10] While Assumption 2 is not testable, this bias makes clear that the problem will be especially problematic if Z is not a strong predictor of FC. A related issue is what is known as the "weak instruments problem". When the instrument and the treatment are weakly associated with one

---

[9] Consider equation (2b), but in a multivariate regression that includes control variables. The variation in foster care that is not related to the controls—variation said to be orthogonal to the controls—is used to estimate β_1. In the 2SLS setting with controls, variation in Z that is not related to the control variables is used to estimate β_1, and the key assumption is that this variation is also unrelated to ε.

[10] In this setting, $\beta_1^{IV} = \text{Cov}(Z,Y)/\text{Cov}(Z,X) = \beta_1 + \text{Cov}(Z,\varepsilon)/\text{Cov}(Z,FC)$.

another, the instrumental-variables estimate is biased toward the usual OLS estimate described in equation (1b). A rule of thumb to avoid this problem is an F-statistic greater than 10 for the F-test that the coefficient(s) on the instrument(s) in the first stage are equal to zero (Stock, Wright, and Yogo, 2002).

2.G. Interpretation

A main concern with the above setup is that the effect of the treatment will vary by person—some children are helped by foster care while others may have worse outcomes. As described in section 2.A., estimates of β_1 aim to estimate an average effect. To the extent that the variation in the instrument only affects a subset of people, then the average will be for those people rather than the population at large. This has been described as a local average treatment effect (Imbens and Angrist, 1994; Dinardo and Lee, 2010). In the example used here, the local average treatment effect is the average effect of foster care placement for those who were placed in care when assigned to a "strict" investigator, but would not have been placed in foster care if they had been assigned to a "lenient" one. These are individuals on the margin of placement—where investigators may disagree about the placement recommendation.

A criticism of the consideration of natural experiments is that these settings may be special in that they do not reveal estimates that are generally applicable. Natural experiments may provide "internal validity"—settings where the instrumental variable assumptions are likely to hold so that observational data may be used to estimate causal effects. The concern is that these estimates may lack "external validity"—the ability to extrapolate these findings to individuals who were not directly affected by the natural experiment.

For instance, two sets of instruments may yield different estimates of β_1—separate estimates for the subgroups affected by these instruments. If these instruments generate variation that are of interest in and of themselves, however, then the local-average treatment effects can aid in interpreting the estimates through the lens provided by the instruments. For example, if the two sets of instruments relate to two different policy reforms, then separate estimates can provide a way to compare the effects of the two reforms. In the current example, the estimates apply to a particularly policy-relevant group: those on the margin of placement.

Further, it is not possible to know exactly which people make up "those that are affected by the natural experiment". Some insights can be gained from the institutional details, and the observable characteristics (e.g. average age), can be described (Abadie, 2003). It can be useful to consider these average characteristics to begin to assess how this treated subgroup compares to the broader population of interest (Doyle 2008, Almond and Doyle, forthcoming).

2.H. Summary

Instrumental-variable strategies exploit variation in treatment that stems from factors that should be unrelated to the individuals being studied. In a social experiment, that factor would be whether the individual were assigned to the treatment group or to the control group. Other times, naturally-occurring randomization due to policy variation can be used to mimic such a trial. A nice way to summarize this strategy is in comparison to using a variable as a control (such as in equation 2b). When a variable is used as a control, the variation in FC that is unrelated to the control variable is used to estimate β_1. The opposite is the case when the variable is used as an instrument: the variation in FC that stems from the instrumental variable is used to estimate β_1.

3. Background on the "Natural Experiment"

Investigators can affect foster care placement in at least three ways. First, the investigator may decide that the case is unsubstantiated, in which case the investigation is unlikely to proceed any further. Second, in emergency situations, the investigator may arrange to have a child removed from home immediately. Third, most foster care placements follow a court hearing by a child-protection judge. The investigator presents evidence of the abuse or neglect to this judge. The investigator can affect the outcome of this hearing by the quality of the investigation conducted and the persuasiveness of the recommendation.

The basis for the natural experiment mimicking random assignment in this context is that the investigators are assigned on a rotational basis with teams. In conversations with case workers, they refer to the case assignment process as "the rotation", and they suggest that it is largely followed in an effort to smooth the work distribution.

Part of the research revealed exceptions that need to be addressed in the estimation. One exception is for cases that are re-investigated—an attempt is made to assign the original investigator. The investigator placement tendency for the child's initial investigator will be the focus of the instrumental-variable calculation. Another exception is that sexual abuse cases and cases of drug-exposed infants are investigated by particular investigators and do not enter into the rotation, and these cases will be excluded from the analysis.

Two other exceptions relate to the group of investigators that a child might be assigned. An attempt is made to assign cases involving Spanish-speaking families to Spanish-speaking investigators, and in some of the areas investigators were assigned to particular neighborhoods. As a result, the rotational assignment is most likely to occur within cells defined by the field

team x neighborhood (ZIP) x Hispanic x year. The Hispanic indicator is part of the cell

definition because it is the best-available proxy for Spanish-speaking cases. These data cells will

be labeled sub-teams for the remainder of the paper.

The proposed instrument uses the foster care placement rate of the initial investigator

assigned to a child's case minus the placement rate of other cases in the same cell. That is,

relative to other investigators that a child could receive, did she receive one that has a high or a

low placement rate? In particular, for child i's initial investigator (the child-protection case

manager) c in sub-team j:

$$(7)\ Z_{c,-j} = \frac{1}{n_{c,-j}} \sum_{k=-j} n_{c,k}\ (\overline{R^{c,k}} - \overline{R^k})$$

where $n_{c,-j}$ is the total number of children investigated by case manager c outside of the

family's sub-team, $n_{c,k}$ is the number of children investigated by case manager c in sub-team k,

$\overline{R^{c,k}}$ is the fraction of children investigated by case manager c in sub-team k that are eventually

removed from home, and $\overline{R^k}$ is the fraction of investigated children in sub-team k that are

eventually removed from home. For ease of exposition, this has been discretized into high vs.

low based on the median of the variable so that the Wald estimator and local-average treatment

effects interpretation are more easily described. The appendix provides more detail on the

continuous version of this instrument.

4. The Data

A unique dataset that links individuals across a wide array of administrative agencies in

Illinois is used to carry out the analysis as described in detail in Doyle (2007). These data are

collected by the Chapin Hall Center for Children, a research institute located at the University of

15

Chicago, and linked together to create the Illinois Integrated Database (Goerge, Van Voorhis, and Lee, 1994). The core of the data comes from the Illinois Department of Children and Family Services, which includes child protection investigation and foster care administrative data. All children investigated for abuse or neglect between July 1, 1990 and December 31, 2000 form the starting point for the analysis. These data have first been linked to the Department of Human Services, where more identifying variables are available. Then they have been linked to the "Delinquency File" of the Juvenile Court of Cook County and to Medicaid Paid Claims data. These data track children who enter the juvenile courts over the same time period. In particular, an appearance before the juvenile court system usually entails three juvenile arrests (or an arrest for a serious charge). This implies that a court appearance identifies a child who has a number of episodes with police and serves as a measure of delinquency. The Medicaid extract includes all fractures, burns, and injuries due to external causes based on the International Classification of Disease, Version 9. Emergencies such as broken bones have the useful property that there is less of a decision to seek care compared to a more chronic condition. Additional healthcare may also be a positive outcome for children, but an increase in these emergency conditions is somewhat easier to interpret as a negative outcome.

Sample Construction

The data are restricted to a subset where the natural experiment is expected to be valid, and where the data are available to match with long-term outcomes. The instrumental variables assumptions described above are assumed to be valid on this sub-sample. Extrapolating the effects to individuals outside of this sample may be unwarranted. This is an example of the tradeoff between internal and external validity described in section 2.

Specifically, the sample is restricted to (1) non-sexual abuse cases (which represent 8% of cases) and non-drug exposure cases (representing another 5% of cases) because these allegations lead to investigations outside of the usual rotational assignment; (2) all children receiving Medicaid prior to the abuse/neglect report to compare children with the same known identifiers. Of the children placed in foster care in Illinois, 82% had received Medicaid prior to the abuse report; and, (3) children who are at least 15 years old in 2000 to observe longer-term outcomes associated with late adolescence. The analysis will focus, then, on children roughly between the ages of 5 and 15 at the time of the abuse investigation. Approximately one-half of children investigated for abuse or neglect are at least 5 years old.

The delinquency outcome necessarily relates to children in Cook County. Another 1% of the observations had missing child characteristics or had too few case manager investigations to calculate the instrument defined below. In a few cases the child was delinquent prior to the abuse report, and these cases are excluded from the delinquency analysis. Last, the analysis is restricted to investigators with at least 10 cases to provide estimates of the placement propensity.[11]

These restrictions result in 15,681 children. In the delinquency sample there are 409 investigators considered, with an average of 38 investigations per case manager used in constructing the measure. In all of the estimates presented, the standard errors reported are clustered at the investigator level. This reflects that the true variation used to estimate the main parameter of interest stems from the investigators.

---

[11] A discussion of the adequacy of 10 cases in estimating the case manager tendencies is discussed in Doyle (2007).

An additional consideration is that for a small number of cases the delinquency occurred between the time of the investigation and the placement, the wait for removal may have contributed to the outcome. It is important not to associate these delinquencies or births with foster care placement. In the outcome comparisons, the indicator for placement is set to zero for these cases.

For ease of comparison, the analysis of emergency healthcare usage employs the same dataset as the one considered for juvenile delinquency in Cook County with an additional restriction. An issue that is especially salient when considering short-term outcomes is that the outcome may happen after the abuse report and precipitate the foster care placement. We would not want to attribute such an outcome to foster care (it may be the cause rather than the effect). A cleaner comparison, at some expense of being less representative of the population at risk of foster care, considers children who were placed in foster care within a month of the initial report. Children who were placed in foster care after the 30 day period are excluded from the analysis, so the comparison is between those placed quickly and those investigated but never placed.

As noted in section 2, the aim of the restrictions is to increase in the internal validity of the estimate—the likelihood that the instrumental-variable assumptions hold. This is likely at the expense of external validity—the effects for children who are placed after many investigations may differ from those who are placed quickly.

The short-term outcome is whether the child first received emergency healthcare between 90 and 365 days from the initial report. By considering emergencies from 90-365 days after the report, and foster care placement within 30 days, the analysis attempts to consider new emergencies as opposed to those caused by the abuse or neglect that led to the initial report.

5. Results

5.A. Juvenile Delinquency Example

Table 1 offers a first look at the plausibility that the rotational assignment is actually used in practice. Recall that the investigator placement differential (7) has been discretized, and it has a mean close to zero and a standard deviation of close to 10%. Mean case characteristics are compared across high- and low-placement propensity investigators. These characteristics are remarkably similar across the two types of investigators. As noted in section two, it is not possible to test whether the instrument is related to the unobserved characteristics (Assumption 2). That said, if high-placement rate investigators had cases that looked like more serious abuse or neglect cases, it would suggest that particular types of cases are assigned to particular investigators and suggest that there is no natural experiment to consider. This does not appear to be the case. Many of these variables are strong predictors of foster care placement, especially the type of reporter and allegation, yet they are not related to the investigator placement tendency.

The treatment of interest for this longer-term outcome is whether the child was ever placed in foster care following his first investigation. Table 2 presents the components of the Wald estimator described by equation (5) for the binary instrument of whether the child's investigator had a relatively high or low placement propensity. These comparisons do not use control variables, which will be discussed below.

Panel A describes the simple first stage. If a child is assigned to an investigator with an above-median placement rate, he has a 9 percentage-point higher likelihood of being placed as

well. This is large relative to a sample placement rate of 22% for those assigned to investigators with a below-median placement rate. This difference is also highly statistically significant.

Panel B reports the reduced form. Children assigned to investigators with a relatively high placement rate are also more likely to be found in the juvenile justice system later in life. The difference is 1.3 percentage points, or 8% higher than the children assigned to investigators with a relatively low placement rate.

Panel C puts these two estimates together. The variation in juvenile delinquency associated with the investigator assignment, divided by the variation in foster care placement associated with this assignment, yields an instrumental-variable estimate of β_1 of 0.14. This can be interpreted as foster care causing a 14 percentage-point increase in juvenile delinquency, at least among marginal cases where the investigator assignment matters.

The implied difference in the IV point estimate is quite large, which suggests some caution in the interpretation. As described by the Wald estimator, the IV estimate stems from an approximately 10 percentage-point increase in the placement probability, which means that the reduced-form change in delinquency is divided by approximately 0.1. If the reduced form is incorrect, the IV estimate would magnify the bias.

That said, the IV estimate is not outside the range of experience among foster children. For example, Courtney, Terao, and Bost (2004) surveyed children who will turn 18 in foster care—those coming from families that could not be reunified and adoption was not completed— and found that two-thirds of the boys and half of the girls had a history of delinquency. In addition, recall that the estimates apply to marginal cases. To the extent that cases where investigators might disagree are less extreme, children who remain home may be much less

likely to suffer poor outcomes relative to those who are placed. In particular, consider if 20% of these marginal cases were placed in foster care, then a 14 percentage-point difference would imply that foster children have a delinquency rate two times higher than that of children who were not placed in foster care.[12]

Table 3 shows results for the first stage, reduced form, and instrumental variables estimates from models similar to equations (3), (4), and (6), respectively.   Estimates are shown with and without control variables, which include indicators for the type of initial reporter, year of age, sex, race, type of allegation, and ZIP code of residence.

Column (1) replicates the first-stage result from Table 2. This estimate in a model with controls is nearly identical: a 9.1 percentage-point increase in the likelihood of foster care placement for those assigned to high-placement-propensity investigators. The standard errors are relatively small at 0.01. Note that the F-statistics are well above 10: over 100 in the model with controls.

Columns (3) and (4) report the reduced form, and, again, the estimates are similar with and without controls. The estimate with controls increases to 0.017 with a standard error of 0.006. This is an 11% higher delinquency rate for those assigned to relatively low-placement rate investigators.

Last, Columns (5) and (6) report the instrumental-variable estimates estimated by 2SLS. As expected from the above results, they are similar with and without controls. With controls,

---

[12] The delinquency rate in the entire sample is 16.4%. This calculation uses the weighted average: 0.8*(JD Risk) + 0.2*(JD Risk + 14.2) = 16.4%. Solving implies the underlying JD risk is 13.5% and foster children are 14.2 percentage points higher at 27.7.

the point estimate increases somewhat from a 14 percentage-point higher delinquency rate due to foster care placement to 18 percentage points.

Overall, the results are similar whether controls are used or not. This is expected given the results of Table 1 and if the randomization has effectively taken place. In that case, the observable and (hopefully) unobservable characteristics should be unrelated to the investigator placement propensity.

5.B. Emergency Healthcare Example

Foster children at the margin of placement may be at higher risk for juvenile delinquency and other risk factors, but that may be counterbalanced by an increase in child safety. A first and admittedly partial look of whether this is the case is offered by the short-term outcome of emergency healthcare. For those who stay at home, emergency healthcare for fractured bones, burns, and other accidents may be related to child abuse or neglect. Meanwhile, for those placed in foster care, previous instrumental-variable results suggest that foster children are more likely to take risks (including the delinquency behavior shown above). This could put foster children at higher risk of such emergencies compared to those who remain at home.

Table 4 reports the results. Columns (1) and (2) show the first stage for this slightly smaller sample and for the 30-day placement indicator. The point estimate of a 7 percentage-point increase in the 30-day foster care placement represents a 56% higher placement rate than the mean of 12.4%.

Columns (3) and (4) show that emergency healthcare within 90-365 days from the report is higher among those who were investigated by relatively high-placement rate investigators. The point estimate is a 1 percentage-point increase, or 11% higher than the mean.

22

Putting these estimates together, Columns (5) and (6) report large increases in emergency care within 90-365 days after the report for those who were placed in foster care within 30 days compared to those who were investigated but never placed. The point estimate suggests a 15 percentage-point higher emergency healthcare rate compared to a mean of 16%. With the lower 30-day placement rate, if 10% of marginal cases are placed in foster care, the point estimates here suggest that emergency care is 3 times higher among those placed in foster care compared to those who were not. It appears that foster care does not reduce such emergencies, but, rather, is associated with a higher incidence.

5.C. Robustness

Doyle (2007) and Doyle (2008) describe a large number of sensitivity analyses. For example, (1) the estimates are qualitatively similar when nonlinear estimation such as a probit or logit models are used; (2) alternative specifications for the instrument, including when the instrument is allowed to be a continuous, rather than a discrete, and when the instrument is calculated using all cases seen by the investigator prior to the child's case; (3) when the estimates focus on different parts of the investigator placement-rate distribution rather than above and below median; (4) across different case types, such as abuse vs. neglect, and boys vs. girls. In addition, clustering the standard errors by the field team rather than the investigator yielded nearly identical standard-error estimates.[13]

---

[13] Table 3 shows that in the model with full controls, the estimated standard error clustered at the investigator level is 0.0632. When the standard errors are clustered at the field team level, the estimated standard error is 0.0634, and when the standard errors are clustered at the sub-team level (defined by field team x ZIP code x Hispanic x year cells), the standard error is 0.0674.

For the emergency healthcare example, the results were similar when the estimation sample was expanded to include the entire state of Illinois. The results were also similar with different definitions of short-term placement such as 10 and 60 days from the report.

6. Interpretation and Conclusions

The results suggest that placing children in foster care increases their likelihood of becoming delinquent during adolescence and requiring emergency healthcare in the short term. Along this one dimension of child safety, it does not appear that foster care is serving a protective role.

The main caveat with these and other instrumental-variable results is that they apply to a particular set of children: those who are affected by the instrument. Consider a child that any investigator and judge would recommend foster care placement. There is simply no variation in the placement decision among these children. Similarly, if no investigator or judge would recommend placement for a particular child, then there is no variation in foster care placement. It is beyond the scope of the empirical strategy to estimate the effects of foster care for these types of children.

The results do apply to a particularly policy-relevant group: those children where the investigator does matter. These are marginal cases where investigators could disagree about how to proceed. This variation is at the heart of the policy question of whether the child-protection system is too aggressive or not aggressive enough.

Further, the results apply to somewhat older children, between the ages of 5 and 15, who were investigated for abuse or neglect in Illinois during the 1990s. To the extent that other foster care systems perform better than this one, the answer could change. Future research that

considers younger children, other states, and other time periods would allow an examination of whether the results apply more generally to child protection policies in the U.S. In order to carry out such an analysis, it is imperative for administrative datasets across a number of programs to be combined into a powerful tool for program evaluation.

References

Abadie, Alberto. (2003) "Semiparametric Instrumental Variable Estimation of Treatment Response Models." *Journal of Econometrics*. 113: 231-263.

Almond, Douglas and Joseph J. Doyle. (forthcoming) "After Midnight: A Regression Discontinuity Design in Length of Postpartum Hospital Stays." *American Economic Journal: Economic Policy*.

Angrist, Joshua D. and Alan B. Krueger. (2001) "Instrumental Variables and the Search for Identification: From Supply and Demand to Natural Experiments" *The Journal of Economic Perspectives*. 15(4): 69-85

Angrist, Joshua D. and Jorn-Steffen Pischke. (2009) *Mostly Harmless Econometrics*. Princeton University Press: Princeton, NJ.

Barth, Richard P. (1999) "After Safety, What is the Goal of Child Welfare Services: Permanency, Family Continuity or Social Benefit?" *International Journal of Social Welfare*. 8: 244-252.

Cameron, Colin A. and Pravin K. Trivedi. (2005) *Microeconometrics: Methods and Applications*. Cambridge University Press: New York. 2005.

Cash, S. J. (2001) "Risk Assessment in Child Welfare: The Art and Science." *Children and Youth Services Review*. 23(11): 811-830.

Clausen, J.M., Landsverk, J., Ganger, W., Chadwick, D., and A. Litrownik. (1998) "Mental Health Problems of Children in Foster Care." *Journal of child and Family Studies*. 7: 283-296.

Courtney, Mark E. (2000) "Research Needed to Improve the Prospects for Children in Out-of-home Placement." *Children and Youth Services Review*. 22:743-761.

Courtney, M.E. and I. Piliavin. (1998) *Foster youths transitions to adulthood: Outcomes 12 to18 months after leaving out-of-home care*. Madison, WI: School of Social Work, University of Wisconsin-Madison. 1998.

Courtney, Mark E., Terao, Sherri, and Noel Bost. (2004) *Midwest Evaluation of the Adult Functioning of Former Foster Youth: Conditions of Youth Preparing to Leave State Care*. Chicago: Chapin Hall Center for Children at the University of Chicago.

Davidson-Arad, Bilha, Englechin-Segal, Dorit, and Yochanan Wozner. (2003) "Short-term follow-up of children at risk: comparison of the quality of life of children removed from home and children remaining at home." *Child Abuse & Neglect.*   27: 733–750.

Dinardo, John and David S. Lee. "Program Evaluation and Research Designs." (forthcoming) *Handbook of Labor Economics*. Vol. IV. Card, David and Orley Ashenfelter (eds).

Doyle, Joseph J. (2007) "Child Protection and Child Outcomes: Measuring the Effects of Foster Care." *American Economic Review*. 97(5): 1583-1610.

Doyle, Joseph J. (2008) "Child Protection and Adult Crime: Using Investigator Assignment to Estimate Causal Effects of Foster Care." *Journal of Political Economy*. 116(4): 746-770.

Dworsky, A. & Courtney, M.E. (2000) *Self-sufficiency of former foster youth in Wisconsin: Analysis of unemployment insurance wage data and public assistance data.* USDHHS: Washington, D.C. 2000.

Elmer, E. (1986) "Outcomes of Residential Treatment for Abused and High-risk Infants." *Child Abuse and Neglect*. 10: 351-360.

Gelles, R.J. (2000) "How Evaluation Research Can Help Reform and Improve the Child Welfare System." *Journal of Aggression, Maltreatment, and Trauma*. 4: 7-28.

Goerge, R. M., Wulczyn, F., and D. Fanshel. (1994) "A Foster Care Research Agenda for the 90s." *Child Welfare*. 73(5): 525-549.

Goerge, R. M., Van Voorhis, J. and B.J. Lee. (1994) "Illinois's longitudinal and relational child and family research database." *Social Science Computer Review*. 12(3): 351-65.

Heckman, James J., LaLonde, Robert, and Jeffrey Smith (1999), "The Economics and Econometrics of Active Labor Market Programs," in O. Ashenfelter and D. Card, Chapter 31, *Handbook of Labor Economics* , Vol. IV, 1865-2073.

Heckman, James J., Ichimura Hidehiko, and Petra Todd (1998), "Matching as an Econometric Evaluation Estimator." *The Review of Economic Studies*. 65: 261-294.

Imbens, Guido W. and Joshua D.Angrist. (1994)   "Identification and Estimation of Local Average Treatment Effects." *Econometrica*. 62(2): 467-475.

James, Sigrid, Landsverk, John, and Donald J. Slymen. (2004) "Placement Movement in Out-of-Home Care: Patterns and Predictors." *Children and Youth Services Review*. 26: 185-206.

Jonson-Reid, M. and R. Barth. (2000) "From Placement to Prison: The path to Adolescent Incarceration from Child Welfare Supervised Foster or Group Care." *Children and Youth Services Review*. 22(7): 493-516.

Kerman, Benjamin, Wildfire, Judith, and Richard P. Barth. (2002) "Outcomes for Young Adults Who Experienced Foster Care.' Children and Youth Services Review, 24(5): 319–344.

Kling, Jeffrey R. 2006. "Incarceration Length, Employment and Earnings." *American Economic Review*, 96(3) (June): 863-876

Lindsey, Duncan. (1994) "Family Preservation and Child Protection: Striking a Balance." *Children and Youth Services Review*. 16(5-6): 279-294.

Maluccio, Anthony N., Pine, Barbara A., and Robin Warsh. (1994) "Protecting Children by Preserving their Families." *Children and Youth Services Review*. 16(5-6): 295-307.

McClellan, Mark., McNeil, Barbara J., and Joseph P. Newhouse. (1994) "Does More Intensive Treatment of Acute Myocardial Infarction in the Elderly Reduce Mortality? Analysis Using Instrumental Variables" *JAMA*. 272(11): 859-866.

McDonald, Thomas P., Allen, Reva I., Westerfelt, Alex, and Irving Piliavin. (1996) *Assessing the Long-Term Effects of Foster Care: A Research Synthesis*. Washington, D.C.: CWLA Press.

National Research Council and Institute of Medicine. (1998) *Violence in Families*. Chalk, Rosemary and Patricia A. King (editors). Washington: National Academy Press.

Newton, Rae, Litronwnik, Alan J., and John A. Landsverk. (2000) "Children and Youth in Foster Care: Disentangling the Relationship Between Problem Behaviors and Number of Placements. *Child Abuse & Neglect*. 24(10): 1363-1374.

Rosenbaum, Paul R. and Donald B. Rubin. (1983) "The Central Role of the Propensity Score in Observational Studies for Causal Effects." *Biometrika*. 70(1): 41-55.

Runyan, D.K., and C. Gould. (1985) "Foster Care for Child Maltreatment: Impact on Delinquent Behavior." *Pediatrics*. 75: 562-68.

Ryan, Joseph P., Hong, Jun Sung, Herz, Denise, and Pedro M. Hernandez. (forthcoming) "Kinship Foster Care and the Risk of Juvenile Delinquency." *Children and Youth Services Review*.

Smith, Dana K., Stormshak, Elizabeth, Chamberlain, Patricia, and Rachel Bridges Whaley. (2001) "Placement Disruption in Treatment Foster Care." *Journal of Emotional and Behavioral Disorders*. 9(3): 200-205.

Stock, James, James Wright, and Motohiro Yogo. (2002). "A Survey of Weak Instruments and Weak Identification in GMM." *Journal of Business and Economic Statistics,* 20(4): 518–29.

Testa, M. & Poertner, J. (In Press). Fostering Accountability: *Using Evidence to Guide and Improve Child Welfare Policy.* Oxford: Oxford University Press.

U.S. Department of Health and Human Services, Administration on Children, Youth and Families. (1999) *Title IV-E Independent Living Programs: a Decade in Review.* Washington, DC: US Government Printing Office. 1999.

US Department of Health and Human Services. (2009) "The AFCARS Report." Washington, DC: US Government Printing Office. October.

Vinnerljung, Bo, Sundell, Knut, Andree Lofholm, Cecilia, and Eva Humlesjo. (2006) "Former Stockholm Child Protection Cases as Young Adults: Do Outcomes Differ Between Those that Received Services and Those that Did Not?" *Children and Youth Services Review*. 28: 59-77.

Wald, M.S., Carlsmith, J.M., and P.H. Leiderman. (1988) *Protecting Abused and Neglected Children*. Stanford, California: Stanford University Press.

Widom, C.S. (1989) The Cycle of Violence, *Science*, 244:160-166.

Wooldridge, Jeffrey, M. (2002) *Econometric Analysis of Cross Section and Panel Data*. Cambridge, MA: MIT Press.

Appendix

The proposed instrument is essentially the placement rate of the child's investigator minus the placement rate of the subteam assigned to the case. This subteam is the group within which the effective randomization is expected to take place. This appendix describes the instrument in a slightly different way to highlight the source of the variation in foster care placement that is used to estimate causal effects.

As in the main text, the structural equation (adding control variables X) for child i is:

$$(A1) \; \text{JD}_i = \beta_0 + \beta_1 \text{FC}_i + \beta_2 X_i + \varepsilon_i$$

To implement the 2-stage Least Squares estimator, the first stage predicts foster care placement:

$$(A2) \; \text{FC}_i = \alpha_0 + \alpha_1 Z_i + \alpha_2 X_i + \mu_i$$

and the second stage substitutes the predicted value from (A2) into (A1) yielding:

$$(A3) \; \text{JD}_i = \beta_0 + \beta_1 \widetilde{FC}_i + \beta_2 X_i + \tilde{\varepsilon}_i$$

The instrument in the main text uses the placement rate of the investigator assigned to child i relative to the other investigators that the child could have been assigned. This was then discretized into an indicator as to whether the relative placement rate of the child's investigator is above or below the median in the sample.

An alternative way to use the within-team difference in placement rates across investigators would be to use investigator fixed effects as the vector of instruments and include

subteam fixed effects to control for differences across teams, denoted δ_t for subteam t that investigates child i: Consider a slightly different structural equation:

$$(A1') \ \text{JD}_i = \beta_0 + \beta_1 \text{FC}_i + \beta_2 X_i + \delta_t + \varepsilon_i$$

δ_t can be thought of as a control that is separated from the X vector.

Now let π_c be a vector of investigator (the case worker) fixed effects for an alternative first stage:

$$(A2') \ \text{FC}_i = \alpha_0 + \alpha_1 X_i + \pi_c + \delta_t + \mu_i$$

And the same second stage, but now the predicted value of FC is calculated using $(A2')$:

$$(A3') \ \text{JD}_i = \beta_0 + \beta_1 \widehat{FC}_i + \beta_2 X_i + \delta_t + \tilde{\varepsilon}_i$$

The inclusion of the subteam fixed effects in $(A2')$ means that the investigator fixed effects will be estimated using within-subteam variation in placement rates across investigators (for more, see Hsiao, 2003). That is, the estimated investigator fixed effects will describe their placement rate relative to others in the subteam. This is analogous to the instrument used in the main text which also used within-subteam variation in placement rates.

The main text uses a summary of these investigator fixed effects as the instrumental variable rather than implement the alternative model described here, although results are similar regardless. The summary measure allows for an easier description of the instrument as a relative placement rate, as well as its relationship with the control variables as in Table 1.

Given that the number of instruments here is the number of investigators, there can be a large number of them. When considering the properties of the estimator as the sample size tends

to infinity, the number of instruments may grow to infinity as well. It has been suggested that such models should be estimated using Limited Information Maximium Likelihood or a jackknife procedure where the individual i is removed from the construction of the fixed effects (Stock, Yogo, Wright, 2002). The main text implements this jackknife procedure by calculating the instrument excluding the child's subteam.

References:

Hsaio, Cheng. (2003) Analysis of Panel Data. Econometric Society Monograph No. 34. Cambridge: Cambridge University Press.

Stock, James, James Wright, and Motohiro Yogo. (2002). "A Survey of Weak Instruments and Weak Identification in GMM." *Journal of Business and Economic Statistics,* 20(4): 518–29.

Table 1: Plausibility

Variable		Investigator Placement Tendency		p-value
		High	**Low**	**p-value**
Initial Reporter	physician	0.128	0.120	0.194
	school	0.127	0.122	0.451
	police	0.134	0.139	0.430
	family	0.296	0.284	0.182
	neighbor	0.052	0.060	0.085
	anonymous	0.140	0.153	0.070
	other government	0.093	0.095	0.774
	other reporter	0.029	0.026	0.418
Age at Report	age	11.4	11.4	0.985
Sex	boy	0.480	0.490	0.237
Race	white	0.109	0.107	0.688
	African American	0.763	0.766	0.687
	Hispanic	0.117	0.118	0.876
Allegation	lack of supervision	0.369	0.369	0.976
	environmental neglect	0.148	0.159	0.128
	other neglect	0.194	0.202	0.340
	abuse	0.437	0.429	0.408
	Observations	15681		

Juvenile Delinquency Sample: Children in Cook County who received an abuse/neglect report between July 1, 1990 & December 31, 2000 and were at least 15 in 2000. p-values calculated using standard errors clustered at the investigator level.

Table 2: Table of Means: Instrumental Variable Estimation

| | | Investigator Placement Propensity | | Difference | p-value |
		High	Low		
A. First Stage	Foster Care Placement	0.316	0.224	0.092	<0.0001
B. Reduced Form	Juvenile Delinquency	0.171	0.158	0.013	0.043
C. IV Estimate	Change in Juvenild Delinquency:	Difference in B ÷ Difference in A:		0.142	
	Change in Foster Care Placement		p-value:	0.035	
	Observations	7792	7889		

Juvenile Delinquency Sample: Children in Cook County who received an abuse/neglect report between July 1, 1990 & December 31, 2000 and were at least 15 in 2000. p-values calculated using standard errors clustered at the investigator level.

Table 3: Introduction of Covariates

	First Stage		Reduced Form		IV Estimates	
Dependent Variable:	Foster Care Placement		Juvenile Delinquency		Juvenile Delinquency	
Estimator:	OLS	OLS	OLS	OLS	2SLS	2SLS
	(1)	(2)	(3)	(4)	(5)	(6)
Investigator Placement Propensity = High	0.092	0.091	0.013	0.017		
	(0.013)**	(0.009)**	(0.0065)**	(0.0058)**		
Foster Care Placement					0.142	0.183
					(0.067)**	(0.063)**
F-statistic (Ho: Above coefficient = 0)	47.8	102				
Full Controls	No	Yes	No	Yes	No	Yes
Mean of Dependent Variable	0.269		0.164		0.164	
Number of Investigators	409					
Observations	15681					

Standard errors are reported, clustered at the investigator level. Full controls include indicators for the type of initial reporter, year of age, sex, race, type of allegation, and ZIP code of residence.

Table 4: Emergency Care within 1 year

Dependent Variable (Time since report):	First Stage <30 Days to FC Placement		Reduced Form Emergency Care within 90-365 Days		IV Estimates Emergency Care within 90-365 Days	
Estimator:	OLS	OLS	OLS	OLS	2SLS	2SLS
	(1)	(2)	(3)	(4)	(5)	(6)
Investigator Placement Propensity = High	0.076	0.071	0.010	0.011		
	(0.0087)**	(0.0073)**	(0.0051)**	(0.0050)**		
Foster Care Placement in < 30 days					0.137	0.156
					(0.066)**	(0.073)**
F-statistic (Ho: Above coefficient = 0)	76.8	94.8				
Full Controls	No	Yes	No	Yes	No	Yes
Mean of Dependent Variable	0.124		0.089		0.089	
Number of Investigators	381					
Observations	12519					

Sample is similar to Table 3, but further restricted to children placed in foster care within 30 days of the abuse or negelect report or
Standard errors are reported, clustered at the case manager level. Full controls include indicators for the type of initial reporter, year
of age, sex, race, type of allegation, and ZIP code of residence.