

What's New in Econometrics

NBER, Summer 2007

Lecture 5, Monday, July 30th, 4.30-5.30pm

Instrumental Variables with Treatment Effect Heterogeneity:

Local Average Treatment Effects

1. INTRODUCTION

Here we investigate the interpretation of instrumental variables estimators allowing for general heterogeneity in the effect of the endogenous regressor. We shall see that instrumental variables estimators generally estimate average treatment effects, with the specific average depending on the choice of instruments. Initially we focus on the case where the endogenous regressor is binary. The example we will use is based on work by Joshua Angrist on estimating the effect of veteran status on earnings (Angrist, 1990). We also discuss the case where the endogenous variable takes on multiple values.

The general theme of this lecture is that with heterogeneous treatment effects, endogeneity creates severe problems for identification of population averages. Population average causal effects are only estimable under very strong assumptions on the effect of the instrument on the endogenous regressor (“identification at infinity”, or under the constant treatment effect assumptions). Without such assumptions we can only identify average effects for subpopulations that are induced by the instrument to change the value of the endogenous regressors. We refer to such subpopulations as *compliers*, and to the average treatment effect that is point identified as the *local average treatment effect*. This terminology stems from the canonical example of a randomized experiment with noncompliance. In this example a random subpopulation is assigned to the treatment, but some of the individuals do not comply with their assigned treatment.

These complier subpopulations are not necessarily the subpopulations that are *ex ante* the most interesting subpopulations, but the data is in general not informative about average effects for other subpopulations without extrapolation, similar to the way in which a randomized experiment conducted on men is not informative about average effects for

women without extrapolation. The set up here allows the researcher to sharply separate the extrapolation to the (sub-)population of interest from exploration of the information in the data. The latter relies primarily on relatively interpretable, and substantively meaningful assumptions and avoids functional form or distributional assumptions. Given estimates for the compliers, one can then use the data to assess the plausibility of extrapolating the local average treatment effect to other subpopulations, using the information on outcomes given one of the two treatment levels and covariates.

With multiple instruments and or with covariates one can assess the evidence for heterogeneity, and the plausibility of extrapolation to the full population more extensively.

2. LINEAR INSTRUMENTAL VARIABLES WITH CONSTANT COEFFICIENTS

First let us briefly review standard linear instrumental variables methods. In the example we are interested in the causal effect of military service on earnings. Let Y_i be the outcome of interest for unit i , W_i the endogenous regressor, and Z_i the instrument. The standard set up is as follows. A linear model is postulated for the relation between the outcome and the endogenous regressor:

$$Y_i = \beta_0 + \beta_1 \cdot W_i + \varepsilon_i.$$

This is a structural/behavioral/causal relationship. There is concern that the regressor W_i is endogenous, that is, that W_i is correlated with ε_i . Suppose that we are confident that a second variable, the instrument Z_i is both uncorrelated with the unobserved component ε_i and correlated with the endogenous regressor W_i . The solution is to use Z_i as an instrument for W_i . There are a couple of ways to implement this.

In Two-Stage-Least-Squares we first estimate a linear regression of the endogenous regressor on the instrument by least squares. Let the estimated regression function be

$$\hat{W}_i = \hat{\pi}_0 + \hat{\pi}_1 \cdot Z_i.$$

Then we regress the outcome on the predicted value of the endogenous regressor, using least

squares:

$$Y_i = \hat{\beta}_0 + \hat{\beta}_1 \cdot \hat{W}_i.$$

Alternatively, with a single instrument we can estimate the two reduced form regressions

$$Y_i = \gamma_0 + \gamma_1 \cdot Z_i + \eta_i, \quad \text{and} \quad W_i = \pi_0 + \pi_1 \cdot Z_i + \nu_i,$$

by least squares and estimate β_1 through Indirect Least Squares (ILS) as the ratio

$$\hat{\beta}_1^{\text{IV}} = \hat{\gamma}_1 / \hat{\pi}_1.$$

If there is a single instrument and single endogenous regressor, we end up in both cases with the ratio of the sample covariance of Y and Z to the sample covariance of W and Z .

$$\hat{\beta}_1^{\text{IV}} = \frac{\frac{1}{N} \sum_{i=1}^N (Y_i - \bar{Y}) \cdot (Z_i - \bar{Z})}{\frac{1}{N} \sum_{i=1}^N (W_i - \bar{W}) \cdot (Z_i - \bar{Z})}.$$

Using a central limit theorem for all the moments and the delta method we can infer the large sample distribution without additional assumptions.

3. POTENTIAL OUTCOMES

First we set up the problem in a slightly different way, using potential outcomes. Let $Y_i(0)$ and $Y_i(1)$ be two potential outcomes for unit i , one for each value of the endogenous regressor or treatment. The first potential outcome $Y_i(0)$ gives the outcome if person i were not to serve in the military, irrespective of whether this person served or not. The second gives the potential outcome given military service, again irrespective of whether the person served or not. We are interested in the causal effect of military service, $Y_i(1) - Y_i(0)$. We cannot directly observe this since we can only observe either $Y_i(0)$ or $Y_i(1)$, but not both. Let W_i be the realized value of the endogenous regressor, equal to zero or one. We observe W_i and

$$Y_i = Y_i(W_i) = \begin{cases} Y_i(1) & \text{if } W_i = 1 \\ Y_i(0) & \text{if } W_i = 0. \end{cases}$$

Now we introduce the instrumental variable set up by defining similar potential outcomes for the treatment. We focus on the case with a binary instrument Z_i . In the Angrist example, Z_i is a binary indicator for having a low draft number, and thus for being draft eligible. Define two potential outcomes $W_i(0)$ and $W_i(1)$, representing the value of the endogenous regressor given the two values for the instrument. The actual or realized value of the endogenous variable is

$$W_i = Y_i(Z_i) = \begin{cases} W_i(1) & \text{if } Z_i = 1 \\ W_i(0) & \text{if } Z_i = 0. \end{cases}$$

So we observe the triple Z_i , $W_i = W_i(Z_i)$ and $Y_i = Y_i(W_i(Z_i))$.

4. LOCAL AVERAGE TREATMENT EFFECTS

4.1. ASSUMPTIONS

The key instrumental variables assumption is

Assumption 1 (INDEPENDENCE)

$$Z_i \perp\!\!\!\perp (Y_i(0), Y_i(1), W_i(0), W_i(1)).$$

It requires that the instrument is as good as randomly assigned, and that it does not directly affect the outcome. The assumption is formulated in a nonparametric way, without definitions of residuals that are tied to functional forms.

It is important to note that this assumption is *not* implied by random assignment of Z_i . To see this, an alternative formulation of the assumption, generalizing the notation slightly, is useful. First we postulate the existence of four potential outcomes, $Y_i(z, w)$, corresponding to the outcome that would be observed if the instrument was $Z_i = z$ and the treatment was $W_i = w$. Then the independence assumption is the combination of two assumptions,

Assumption 2 (RANDOM ASSIGNMENT)

$$Z_i \perp\!\!\!\perp (Y_i(0, 0), Y_i(0, 1), Y_i(1, 0), Y_i(1, 1), W_i(0), W_i(1)).$$

and

Assumption 3 (EXCLUSION RESTRICTION)

$$Y_i(z, w) = Y_i(z', w), \quad \text{for all } z, z', w.$$

The first of these two assumptions is implied by random assignment of Z_i , but the second is substantive, and randomization has no bearing on it.

It is useful for our approach to think about the compliance behavior of the different units, that is how they respond to different values of the instrument in terms of the treatment received. Table 1 gives the four possible pairs of values $(W_i(0), W_i(1))$, given the binary nature of the treatment and instrument: We cannot directly establish the type of a unit based

Table 1: COMPLIANCE TYPES

		$W_i(0)$	
		0	1
$W_i(1)$	0	never-taker	defier
	1	complier	always-taker

on what we observe for them since we only see the pair (Z_i, W_i) , not the pair $(W_i(0), W_i(1))$. Nevertheless, we can rule out some possibilities. Table 2 summarizes the information about compliance behavior from observed treatment status and instrument.

To make additional progress we we consider a *monotonicity* assumption, also known as the *no-defiers* assumption:

Assumption 4 (MONOTONICITY/NO-DEFIERS)

$$W_i(1) \geq W_i(0).$$

Table 2: COMPLIANCE TYPE BY TREATMENT AND INSTRUMENT

		Z_i	
		0	1
W_i	0	complier/never-taker	never-taker/defier
	1	always-taker/defier	complier/always-taker

This assumption makes sense in a lot of applications. It is implied directly by many (constant coefficient) latent index models of the type:

$$W_i(z) = 1\{\pi_0 + \pi_1 \cdot z + \varepsilon_i > 0\},$$

but it is much weaker than that. For example, one can allow for π_1 to vary across the population, as long as it is the same sign for all units. In the canonical non-compliance example this assumption is very plausible: if Z_i is assignment to a treatment, and W_i is an indicator for receipt of treatment, it makes sense that there are few, if any, individuals who always to the exact opposite of what their assignment is.

4.2. THE LOCAL AVERAGE TREATMENT EFFECT

Given this monotonicity assumption the information we can extract from observed compliance behavior increases.

Table 3: COMPLIANCE TYPE BY TREATMENT AND INSTRUMENT GIVEN MONOTONICITY

		Z_i	
		0	1
W_i	0	complier/never-taker	never-taker
	1	always-taker	complier/always-taker

Let π_c , π_n , and π_a be the population proportions of compliers, never-takers and always-takers respectively. We can estimate those from the population distribution of treatment and instrument status:

$$\mathbb{E}[W_i|Z_i = 0] = \pi_a, \quad \mathbb{E}[W_i|Z_i = 1] = \pi_a + \pi_c,$$

which we can invert to infer the population shares of the different types:

$$\pi_a = \mathbb{E}[W_i|Z_i = 0], \quad \pi_c = \mathbb{E}[W_i|Z_i = 1] - \mathbb{E}[W_i|Z_i = 0],$$

and

$$\pi_n = 1 - \mathbb{E}[W_i|Z_i = 1].$$

Now consider average outcomes by instrument and treatment status:

$$\mathbb{E}[Y_i|W_i = 0, Z_i = 0] = \frac{\pi_c}{\pi_c + \pi_n} \cdot \mathbb{E}[Y_i(0)|\text{complier}] + \frac{\pi_n}{\pi_c + \pi_n} \cdot \mathbb{E}[Y_i(0)|\text{never-taker}],$$

$$\mathbb{E}[Y_i|W_i = 0, Z_i = 1] = \mathbb{E}[Y_i(0)|\text{never-taker}],$$

$$\mathbb{E}[Y_i|W_i = 1, Z_i = 0] = \mathbb{E}[Y_i(1)|\text{always-taker}],$$

and

$$\mathbb{E}[Y_i|W_i = 1, Z_i = 1] = \frac{\pi_c}{\pi_c + \pi_a} \cdot \mathbb{E}[Y_i(1)|\text{complier}] + \frac{\pi_a}{\pi_c + \pi_a} \cdot \mathbb{E}[Y_i(1)|\text{always-taker}].$$

From these relationships we can infer the average outcome by treatment status for compliers,

$$\mathbb{E}[Y_i(0)|\text{complier}], \quad \text{and} \quad \mathbb{E}[Y_i(1)|\text{complier}],$$

and thus the average effect for compliers:

$$\mathbb{E}[Y(1) - Y_i(0)|\text{complier}] = \mathbb{E}[Y_i(1)|\text{complier}] - \mathbb{E}[Y_i(0)|\text{complier}].$$

We can also get there another way. Consider the least squares regression of Y on a constant and Z . The slope coefficient in that regression estimates

$$\mathbb{E}[Y_i|Z_i = 1] - \mathbb{E}[Y_i|Z_i = 0].$$

Consider the first term:

$$\begin{aligned} \mathbb{E}[Y_i|Z_i = 1] &= \mathbb{E}[Y_i|Z_i = 1, \text{complier}] \cdot \Pr(\text{complier}|Z_i = 1) \\ &\quad + \mathbb{E}[Y_i|Z_i = 1, \text{never - taker}] \cdot \Pr(\text{never - taker}|Z_i = 1) \\ &\quad + \mathbb{E}[Y_i|Z_i = 1, \text{always - taker}] \cdot \Pr(\text{always - taker}|Z_i = 1) \\ &= \mathbb{E}[Y_i(1)|\text{complier}] \cdot \pi_c \\ &\quad + \mathbb{E}[Y_i(0)|\text{never - taker}] \cdot \pi_0 + \mathbb{E}[Y_i(1)|\text{always - taker}] \cdot \pi_a. \end{aligned}$$

Similarly

$$\begin{aligned} \mathbb{E}[Y_i|Z_i = 0] &= \mathbb{E}[Y_i|Z_i = 0, \text{complier}] \cdot \Pr(\text{complier}|Z_i = 0) \\ &\quad + \mathbb{E}[Y_i|Z_i = 0, \text{never - taker}] \cdot \Pr(\text{never - taker}|Z_i = 0) \\ &\quad + \mathbb{E}[Y_i|Z_i = 0, \text{always - taker}] \cdot \Pr(\text{always - taker}|Z_i = 0) \\ &= \mathbb{E}[Y_i(0)|\text{complier}] \cdot \pi_c \\ &\quad + \mathbb{E}[Y_i(0)|\text{never - taker}] \cdot \pi_0 + \mathbb{E}[Y_i(1)|\text{always - taker}] \cdot \pi_a. \end{aligned}$$

Hence the difference is

$$\mathbb{E}[Y_i|Z_i = 1] - \mathbb{E}[Y_i|Z_i = 0] = \mathbb{E}[Y_i(1) - Y_i(0)|\text{complier}] \cdot \pi_c.$$

The same argument can be used to show that the slope coefficient in the regression of W on Z is

$$\mathbb{E}[W_i|Z_i = 1] - \mathbb{E}[W_i|Z_i = 0] = \pi_c.$$

Hence the instrumental variables estimand, the ratio of these two reduced form estimands, is equal to the local average treatment effect

$$\beta^{\text{IV}} = \frac{\mathbb{E}[Y_i|Z_i = 1] - \mathbb{E}[Y_i|Z_i = 0]}{\mathbb{E}[W_i|Z_i = 1] - \mathbb{E}[W_i|Z_i = 0]} = \mathbb{E}[Y_i(1) - Y_i(0)|\text{complier}].$$

The key insight is that the data are informative solely about the average effect for compliers only. Put differently, the data are not informative about the average effect for nevertakers because they are never seen receiving the treatment, and they are not informative about the average effect for alwaystakers because they are never seen without the treatment. A similar insight in a parametric setting is discussed in Björklund and Moffitt (1987).

A special case of considerable interest is that with one-side non-compliance. Suppose that $W_i(0) = 0$, so that those assigned to the control group cannot receive the active treatment (but those assigned to the active treatment can decline to take it). In that case only two compliance types remain, compliers and always-takers. Monotonicity is automatically satisfied. The average effect for compliers is now equal to the average effect for the treated, since any one receiving the treatment is by definition a complier. This case was first studied in Bloom (1984).

4.3 EXTRAPOLATING TO THE FULL POPULATION

Although we cannot consistently estimate the average effect of the treatment for always-takers and never-takers, we do have some information about the outcomes for these subpopulations given one of the two treatment arms. Specifically, we can estimate

$$\mathbb{E}[Y_i(0)|\text{never-taker}], \quad \text{and} \quad \mathbb{E}[Y_i(1)|\text{always-taker}].$$

We can learn from these averages whether there is any evidence of heterogeneity in outcomes by compliance status, by comparing the pair of average outcomes of $Y_i(0)$;

$$\mathbb{E}[Y_i(0)|\text{never-taker}], \quad \text{and} \quad \mathbb{E}[Y_i(0)|\text{complier}],$$

and the pair of average outcomes of $Y_i(1)$:

$$\mathbb{E}[Y_i(1)|\text{always-taker}], \quad \text{and} \quad \mathbb{E}[Y_i(1)|\text{complier}].$$

If compliers, never-takers and always-takers are found to be substantially different in levels, by evidence of substantial difference between $\mathbb{E}[Y_i(0)|\text{never-taker}]$ and $\mathbb{E}[Y_i(0)|\text{complier}]$, and or/between $\mathbb{E}[Y_i(1)|\text{always-taker}]$, and $\mathbb{E}[Y_i(1)|\text{complier}]$, then it appears much less plausible that the average effect for compliers is indicative of average effects for other compliance types. On the other hand, if one finds that outcomes given the control treatment for never-takers and compliers are similar, and outcomes given the treatment are similar for compliers and always-takers, it is more plausible that average treatment effects for these groups are also comparable.

4.4 COVARIATES

The local average treatment effect result implies in general that one cannot consistently estimate average effects for subpopulations other than compliers. This still holds in cases where we observe covariates. One can incorporate the covariates into the analysis in a number of different ways. Traditionally the TSLS set up is used with the covariates entering in the outcome equation linearly and additively, as

$$Y_i = \beta_0 + \beta_1 \cdot W_i + \beta_2' X_i + \varepsilon_i,$$

with the covariates added to the set of instruments. Given the potential outcome set up with general heterogeneity in the effects of the treatment, one may also wish to allow for more heterogeneity in the correlations with the covariates. Here we describe a general way of doing so. Unlike TSLS type approaches, this involves modelling both the dependence of

the outcome and the treatment on the covariates. Although there is often a reluctance to model the relation between the treatment, there appears no particular reason that economic theory is more informative about the relation between covariates and outcomes than about the relation between covariates and the choices that lead to the treatment.

A full model can be decomposed into two parts, a model for the compliance type given covariates, and a model for the potential outcomes given covariates for each compliance type. A traditional parametric model with a dummy endogenous variables might have the form (translated to the potential outcome set up used here):

$$W_i(z) = 1\{\pi_0 + \pi_1 \cdot z + \pi_2'X_i + \eta_i \geq 0\},$$

$$Y_i(w) = \beta_0 + \beta_1 \cdot w + \beta_2'X_i + \varepsilon_i,$$

with (η_i, ε_i) jointly normally distributed (e.g., Heckman, 1978). Such a model can be viewed as imposing various restrictions on the relation between compliance types, covariates and outcomes. For example, in this model, if $\pi_1 > 0$, compliance type depends on η_i :

$$\text{unit } i \text{ is a } \begin{cases} \text{never-taker} & \text{if } \eta_i < -\pi_0 - \pi_1 - \pi_2'X_i \\ \text{complier} & \text{if } -\pi_0 - \pi_1 - \pi_2'X_i \leq \eta_i < -\pi_0 - \pi_1 - \pi_2'X_i \\ \text{always-taker} & \text{if } -\pi_0 - \pi_2'X_i \leq \eta_i, \end{cases}$$

which imposes strong restrictions on the relationship between type and outcomes.

An alternative approach is to model the potential outcome $Y_i(w)$ for units with compliance type t given covariates X_i through a common functional form with type and treatment specific parameters:

$$f_{Y(w)|X,T}(y(w)|x, t) = f(y|x; \theta_{wt}),$$

for $(w, t) = (0, n), (0, c), (1, c), (1, a)$. A natural model for the distribution of type is a trinomial logit model:

$$\Pr(T_i = \text{complier}|X_i) = \frac{1}{1 + \exp(\pi_n'X_i) + \exp(\pi_a'X_i)},$$

$$\Pr(T_i = \text{never-taker} | X_i) = \frac{\exp(\pi'_n X_i)}{1 + \exp(\pi'_n X_i) + \exp(\pi'_a X_i)},$$

and

$$\Pr(T_i = \text{always-taker} | X_i) = 1 - \Pr(T_i = \text{complier} | X_i) - \Pr(T_i = \text{never-taker} | X_i).$$

The log likelihood function is then, factored in terms of the contribution by observed W_i, Z_i values:

$$\begin{aligned} \mathcal{L}(\pi_n, \pi_a, \theta_{0n}, \theta_{0c}, \theta_{1c}, \theta_{1a}) = & \\ & \times \prod_{i|W_i=0, Z_i=1} \frac{\exp(\pi'_n X_i)}{1 + \exp(\pi'_n X_i) + \exp(\pi'_a X_i)} \cdot f(Y_i | X_i; \theta_{0n}) \\ & \times \prod_{i|W_i=0, Z_i=0} \left(\frac{\exp(\pi'_n X_i)}{1 + \exp(\pi'_n X_i)} \cdot f(Y_i | X_i; \theta_{0n}) + \frac{1}{1 + \exp(\pi'_n X_i)} \cdot f(Y_i | X_i; \theta_{0c}) \right) \\ & \times \prod_{i|W_i=1, Z_i=1} \left(\frac{\exp(\pi'_a X_i)}{1 + \exp(\pi'_a X_i)} \cdot f(Y_i | X_i; \theta_{1a}) + \frac{1}{1 + \exp(\pi'_a X_i)} \cdot f(Y_i | X_i; \theta_{1c}) \right) \\ & \times \prod_{i|W_i=1, Z_i=0} \frac{\exp(\pi'_a X_i)}{1 + \exp(\pi'_n X_i) + \exp(\pi'_a X_i)} \cdot f(Y_i | X_i; \theta_{1a}). \end{aligned}$$

For example, the second factor consists of the contributions of individuals with $Z_i = 0, W_i = 0$, who are known to be either compliers or never-takers. Maximizing this is straightforward using the EM algorithm (Dempster, Laird, and Rubin, 1977). For an empirical example of this approach see Hirano, Imbens, Rubin and Zhou (2000), and Imbens and Rubin (1997).

In small samples one may wish to incorporate restrictions on the effects of the covariates, and for example assume that the effect of covariates on the outcome is the same irrespective of compliance type. An advantage of this approach is that it can easily be generalized. The type probabilities are nonparametrically identified as functions of the covariates, and the similarly the outcome distributions by type as a function of the covariates.

5. EFFECTS OF MILITARY SERVICE ON EARNINGS

Angrist (1989) was interested in estimating the effect of serving in the military on earnings. Angrist was concerned about the possibility that those choosing to serve in the military are different from those who do not in ways that affects their subsequent earnings irrespective of serving in the military. To avoid biases in simple comparisons of veterans and non-veterans, he exploited the Vietnam era draft lottery. Specifically he uses the binary indicator whether or not your draft lottery number made you eligible to be drafted as an instrument. This was tied to an individual's day of birth, so more or less random. Even so, that does not make it valid as an instrument. As the outcome of interest Angrist uses log earnings.

The simple ols regression leads to:

$$\log(\widehat{\text{earnings}})_i = 5.4364 - 0.0205 \cdot \widehat{\text{veteran}}_i$$

(0079) (0.0167)

In Table 4 we present population sizes of the four treatment/instrument samples. For example, with a low lottery number 5,948 individuals do not, and 1,372 individuals do serve in the military.

Table 4: TREATMENT STATUS BY ASSIGNMENT

		Z_i	
		0	1
W_i	0	5,948	1,915
	1	1,372	865

Using these data we get the following proportions of the various compliance types, given in Table 5, under the non-defiers assumption. For example, the proportion of nevertakers is

estimated as the conditional probability of $W_i = 0$ given $Z_i = 1$:

$$\Pr(\text{nevertaker}) = \frac{1915}{1915 + 865}.$$

Table 5: COMPLIANCE TYPES: ESTIMATED PROPORTIONS

		$W_i(0)$	
		0	1
$W_i(1)$	0	never-taker (0.6888)	defier (0)
	1	complier (0.1237)	always-taker (0.3112)

Table 6 gives the average outcomes for the four groups, by treatment and instrument status.

Table 6: ESTIMATED AVERAGE OUTCOMES BY TREATMENT AND INSTRUMENT

		Z_i	
		0	1
W_i	0	$\widehat{\mathbb{E}[Y]} = 5.4472$	$\widehat{\mathbb{E}[Y]} = 5.4028$
	1	$\widehat{\mathbb{E}[Y]} = 5.4076,$	$\widehat{\mathbb{E}[Y]} = 5.4289$

Table 7 gives the estimated averages for the four compliance types, under the exclusion restriction. This restriction is the key assumption here. There are a number of reasons why it may be violated, e.g., never-takers taking active actions to avoid military service if draft eligible. The local average treatment effect is -0.2336, a 23% drop in earnings as a result of serving in the military.

Simply doing IV or TSLS would give you the same numerical results:

$$\log(\widehat{\text{earnings}})_i = 5.4836 - 0.2336 \cdot \widehat{\text{veteran}}_i$$

Table 7: COMPLIANCE TYPES: ESTIMATED AVERAGE OUTCOMES

		$W_i(0)$	
		0	1
$W_i(1)$	0	never-taker: $\widehat{\mathbb{E}[Y_i(0)]} = 5.4028$	
	1	complier: $\widehat{\mathbb{E}[Y_i(0)]} = 5.6948$, $\widehat{\mathbb{E}[Y_i(1)]} = 5.4612$ always-taker: $\widehat{\mathbb{E}[Y_i(1)]} = 5.4076$	
		(0.0289)	(0.1266)

It is interesting in this application to inspect the average outcome for different compliance groups. Average log earnings for never-takers are 5.40, lower by 29% than average earnings for compliers who do not serve in the military. This suggests that never-takers are substantially different than compliers, and that the average effect of 23% for compliers need not be informative never-takers. In contrast, average log earnings for always-takers are only 6% lower than those for compliers who serve, suggesting that the differences between always-takers and compliers are considerably smaller.

6. MULTIVALUED INSTRUMENTS

For any two values of the instrument z_0 and z_1 satisfying the local average treatment effect assumptions we can define the corresponding local average treatment effect:

$$\tau_{z_1, z_0} = \mathbb{E}[Y_i(1) - Y_i(0) | W_i(z_1) = 1, W_i(z_0) = 0].$$

Note that these local average treatment effects need not be the same for different pairs of instrument values. Comparisons of estimates based on different instruments underlies tests of overidentifying restrictions in TSLS settings. An alternative interpretation of rejections in such testing procedures is therefore that the effects of interest vary, rather than that some of the instruments are invalid. Without assuming homogenous effects there are no tests in

general for the validity of the instruments.

The presence of multi-valued, or similarly, multiple, instruments, does, however, provide an opportunity to assess variation in treatment effects, as well as an opportunity to obtain average effects for subpopulations closer to the one of ultimate interest. Suppose that we have an instrument Z_i with support z_0, z_1, \dots, z_K . Suppose also that the monotonicity assumption holds for all pairs z and z' , and suppose that the instruments are ordered in such a way that

$$p(z_{k-1}) \leq p(z_k), \quad \text{where } p(z) = \mathbb{E}[W_i | Z_i = z].$$

Also suppose that the instrument is relevant,

$$\mathbb{E}[g(Z_i) \cdot W_i] \neq 0.$$

Then the instrumental variables estimator based on using $g(Z)$ as an instrument for W estimates a weighted average of local average treatment effects:

$$\tau_{g(\cdot)} = \frac{\text{Cov}(Y_i, g(Z_i))}{\text{Cov}(W_i, g(Z_i))} = \sum_{k=1}^K \lambda_k \cdot \tau_{z_k, z_{k-1}},$$

where

$$\lambda_k = \frac{(p(z_k) - p(z_{k-1})) \cdot \sum_{l=k}^K \pi_l (g(z_l) - \mathbb{E}[g(Z_i)])}{\sum_{k=1}^K (p(z_k) - p(z_{k-1})) \cdot \sum_{l=k}^K \pi_l (g(z_l) - \mathbb{E}[g(Z_i)])},$$

$$\pi_k = \Pr(Z_i = z_k).$$

These weights are nonnegative and sum up to one.

By choosing $g(z)$ one can choose a different weight function, although there is obviously a limit to what one can do. One can only estimate a weighted average of the local average treatment effects defined for all pairs of instrument values in the support of the instrument.

If the instrument is continuous, and $p(z)$ is continuous in z , we can define the limit of the local average treatment effects

$$\tau_z = \lim_{z' \downarrow z_0, z'' \uparrow z_0} \tau_{z', z''}.$$

In this case with the monotonicity assumption hold for all pairs z and z' , we can use the implied structure on the compliance behavior by modelling $W_i(z)$ as a threshold crossing process,

$$W_i(z) = 1\{h(z) + \eta_i \geq 0\},$$

with the scalar unobserved component η_i independent of the instrument Z_i . This type of latent index model is used extensively in work by Heckman (Heckman and Robb, 1985; Heckman, 1990; Heckman and Vytlacil, 2005), as well as in Vytlacil (2000). Vytlacil shows that if the earlier three assumptions hold for all pairs z and z' , then there is a function $h(\cdot)$ such that this latent index structure is consistent with the joint distribution of the observables. The latent index structure implies that individuals can be ranked in terms of an unobserved component η_i such that if for two individuals i and j we have $\eta_i > \eta_j$, then $W_i(z) \geq W_j(z)$ for all z .

Given this assumption, we can define the marginal treatment effect $\tau(\eta)$ as

$$\tau(\eta) = \mathbb{E}[Y_i(1) - Y_i(0) | \eta_i = \eta].$$

This marginal treatment effect relates directly to the limit of the local average treatment effects

$$\tau(\eta) = \tau_z, \quad \text{with } \eta = -h(z).$$

Note that we can only define this for values of η for which there is a z such that $\tau = -h(z)$. Normalizing the marginal distribution of η to be uniform on $[0, 1]$ (Vytlacil, 2002), this

restricts η to be in the interval $[\inf_z p(z), \sup_z p(z)]$, where $p(z) = \Pr(W_i = 1 | Z_i = z)$. Heckman and Vytlacil (2005) characterize various average treatment effects in terms of this limit. For example, the average treatment effect is simply the average of the marginal treatment effect over the marginal distribution of η :

$$\tau = \int_{\eta} \tau(\eta) dF_{\eta}(\eta).$$

In practice the same limits remain on the identification of average effects. The population average effect is only identified if the instrument moves the probability of participation from zero to one. In fact identification of the population average treatment effect does not require identification of $\tau(\eta)$ at every value of η . The latter is sufficient, but not necessary. For example, in a randomized experiment (corresponding to a binary instrument with the treatment indicator equal to the instrument) the average treatment effect is obviously identified, but the marginal treatment effect is not for any value of η .

7. MULTIVALUED ENDOGENOUS VARIABLES

Now suppose that the endogenous variable W takes on values $0, 1, \dots, J$. We still assume that the instrument Z is binary. We study the interpretation of the instrumental variables estimand

$$\tau = \frac{\text{Cov}(Y_i, Z_i)}{\text{Cov}(W_i, Z_i)} = \frac{\mathbb{E}[Y_i | Z_i = 1] - \mathbb{E}[Y_i | Z_i = 0]}{\mathbb{E}[W_i | Z_i = 1] - \mathbb{E}[W_i | Z_i = 0]}.$$

We make the exclusion assumption that

$$Y_i(w) \perp W_i(z) \perp Z_i,$$

and a version of the monotonicity assumption,

$$W_i(1) \geq W_i(0),$$

Then we can write the instrumental variables estimand as

$$\tau = \sum_{j=1}^J \lambda_j \cdot \mathbb{E}[Y_i(j) - Y_i(j-1) | W_i(1) \geq j > W_i(0)],$$

where

$$\lambda_j = \frac{\Pr(W_i(1) \geq j > W_i(0))}{\sum_{i=1}^J \Pr(W_i(1) \geq i > W_i(0))}.$$

Note that we can estimate the weights λ_j because

$$\begin{aligned} \Pr(W_i(1) \geq j > W_i(0)) &= \Pr(W_i(1) \geq j) - \Pr(W_i(0) \geq j) \\ &= \Pr(W_i(1) \geq j | Z_i = 1) - \Pr(W_i(0) \geq j | Z_i = 0) \\ &= \Pr(W_i \geq j | Z_i = 1) - \Pr(W_i \geq j | Z_i = 0), \end{aligned}$$

using the monotonicity assumption.

8. INSTRUMENTAL VARIABLES ESTIMATES OF THE RETURNS TO EDUCATION USING QUARTER OF BIRTH AS AN INSTRUMENT

Here we use a subset of the data used by Angrist and Krueger in their 1991 study of the returns to education. Angrist and Krueger were concerned with the endogeneity of education, worrying that individuals with higher ability would have had higher earnings given any level of education, as well as be more likely to have high levels of education. In that case simple least squares estimates would over estimate the returns to education. Their idea was that individuals born in different parts of the year are subject to slightly different compulsory schooling laws. If you are born before a fixed cutoff date you enter school at a younger age than if you are born after that cutoff date, and given that you are allowed to leave school when you turn sixteen, those individuals born before the cutoff date are required to complete more years of schooling. The instrument can therefore be thought of

as the tightness of the compulsory schooling laws, with the tightness being measured by the individual's quarter of birth.

Angrist and Krueger implement this using census data with quarter of birth indicators as the instrument. Table 1 gives average years of education and sample sizes by quarter of birth.

Table 8: AVERAGE LEVEL OF EDUCATION BY QUARTER OF BIRTH

quarter	1	2	3	4
average level of education	12.69	12.74	12.81	12.84
standard error	0.01	0.01	0.01	0.01
number of observations	81,671	80,138	86,856	80,844

In the illustrations below we just use a single instrument, an indicator for being born in the first quarter. First let us look at the reduced form regressions of log earnings and years of education on the first quarter of birth dummy:

$$\widehat{\text{educ}}_i = 12.797 - 0.109 \cdot \text{qob}_i$$

(0.006) (0.013)

and

$$\log(\widehat{\text{earnings}})_i = 5.903 - 0.011 \cdot \text{qob}_i$$

(0.001) (0.003)

The instrumental variables estimate is the ratio of the reduced form coefficients,

$$\hat{\beta}^{\text{IV}} = \frac{-0.1019}{-0.011} = 0.1020.$$

Now let us interpret this in the context of heterogeneous returns to education. This estimate is an average of returns to education, consisting of two types of averaging. The first is over different levels of education. That is, it is a weighted average of the return to moving from nine to ten years, to moving from ten to eleven years, to moving from eleven to twelve years, etcetera. In addition, for any level, e.g., to moving from nine to ten years of education, it is an average effect where the averaging is over those people whose schooling would have been at least ten years of education if tighter compulsory schooling laws had been in effect for them, and who would have had less than ten years of education had they been subject to the looser compulsory schooling laws.

Furthermore, we can estimate how large a fraction of the population is in these categories. First we estimate the

$$\gamma_j = \Pr(W_i(1) \geq j > W_i(0)) = \Pr(W_i \geq j | Z_i = 1) - \Pr(W_i \geq j | Z_i = 0)$$

as

$$\hat{\gamma}_j = \frac{1}{N_1} \sum_{i|Z_i=1} 1\{W_i \geq j\} - \frac{1}{N_0} \sum_{i|Z_i=0} 1\{W_i \geq j\}.$$

This gives the unnormalized weight function. We then normalize the weights so they add up to one, $\hat{\lambda}_j = \hat{\gamma}_j / \sum_i \hat{\gamma}_i$.

Figure 1-4 present some of the relevant evidence here. First, Figure 1 gives the distribution of years of education. Figure 2 gives the normalized and Figure 3 gives the unnormalized weight functions. Figure 4 gives the distribution functions of years of education by the two values of the instrument. The most striking feature of these figures (not entirely unanticipated) is that the proportion of individuals in the “complier” subpopulations is extremely small, never more than 2% of the population. This implies that these instrumental variables estimates are averaged only over a very small subpopulation, and that there is little reason to believe that they generalize to the general population. (Nevertheless, this may well be a very interesting subpopulation for some purposes.) The nature of the instrument also suggests

that most of the weight would be just around the number of years that would be required under the compulsory schooling laws. The weight function is actually much flatter, putting weight even on fourteen to fifteen years of education.

Figure 1: histogram estimate of density of years of education




Figure 2: Normalized Weight Function for Instrumental Variables Estimand




Figure 3: Unnormalized Weight Function for Instrumental Variables Estimand





Figure 3: Education Distribution Function by Quarter



REFERENCES

ANGRIST, J. D., AND G. W. IMBENS, (1995), "Two-Stage Least Squares Estimation of Average Causal Effects in Models with Variable Treatment Intensity," *Journal of the American Statistical Association*, Vol 90, No. 430, 431-442.

ANGRIST, J.D., G.W. IMBENS AND D.B. RUBIN (1996), "Identification of Causal Effects Using Instrumental Variables," (with discussion) *Journal of the American Statistical Association*, 91, 444-472.

ANGRIST, J., (1990), "Lifetime Earnings and the Vietnam Era Draft Lottery: Evidence from Social Security Administrative Records," *American Economic Review*, 80, 313-335.

ANGRIST, J. AND A. KRUEGER, (1992), "The Effect of Age at School Entry on Educational Attainment: An Application of Instrumental Variables with Moments from Two Samples," *Journal of the American Statistical Association* 87, June.

BJÖRKLUND, A. AND R. MOFFITT, (1987), "The Estimation of Wage Gains and Welfare Gains in Self-Selection Models", *Review of Economics and Statistics*, Vol. LXIX, 42-49.

BLOOM, H., (1984), "Accounting for No-shows in Experimental Evaluation Designs," *Evaluation Review*, 8(2) 225-246.

DEMPSTER, A., N. LAIRD, AND D. RUBIN (1977), "Maximum Likelihood Estimation from Incomplete Data Using the EM Algorithm (with discussion)," *Journal of the Royal Statistical Society, Series B*, 39, 1-38.

HECKMAN, J. J. (1990), "Varieties of Selection Bias," *American Economic Review* 80, 313-318.

HECKMAN, J., AND R. ROBB, (1985), "Alternative Methods for Evaluating the Impact of Interventions," in Heckman and Singer (eds.), *Longitudinal Analysis of Labor Market Data*, Cambridge, Cambridge University Press.

HECKMAN, J., AND E. VYTLACIL, (2005), "Structural Equations, Treatment Effects,

and Econometric Policy Evaluation,” *Econometrica*, Vol. 73(3), 669-738.

HIRANO, K., G. IMBENS, D. RUBIN, AND X. ZHOU (2000), “Identification and Estimation of Local Average Treatment Effects,” *Biostatistics*, Vol. 1(1), 69-88.

IMBENS, G., AND J. ANGRIST (1994), “Identification and Estimation of Local Average Treatment Effects,” *Econometrica*, Vol. 61, No. 2, 467-476.

IMBENS, G. W., AND D. B. RUBIN, (1997), “Bayesian Inference for Causal Effects in Randomized Experiments with Noncompliance,” *Annals of Statistics*, Vol. 25, No. 1, 305–327.

VYTLACIL, E., (2002), “Independence, Monotonicity, and Latent Index Models: An Equivalence Result,” *Econometrica*, Vol. 70(1), 331-341.