

NBER WORKING PAPER SERIES

INSTRUMENTAL VARIABLES AND THE SEARCH FOR IDENTIFICATION: FROM
SUPPLY AND DEMAND TO NATURAL EXPERIMENTS

Joshua D. Angrist
Alan B. Krueger

Working Paper 8456
<http://www.nber.org/papers/w8456>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
September 2001

This paper was prepared for the *Journal of Economic Perspectives* symposium on econometric tools and presented at the January 2001 meetings of the American Economic Association. We are grateful to Brad DeLong, David Freedman, Tim Taylor, and Michael Waldman for helpful comments. The views expressed herein are those of the authors and not necessarily those of the National Bureau of Economic Research.

© 2001 by Joshua D. Angrist and Alan B. Krueger. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Instrumental Variables and the Search for Identification:
From Supply and Demand to Natural Experiments
Joshua D. Angrist and Alan B. Krueger
NBER Working Paper No. 8456
September 2001
JEL No. G10, J31, B23

ABSTRACT

The method of instrumental variables was first used in the 1920s to estimate supply and demand elasticities, and later used to correct for measurement error in single-equation models. Recently, instrumental variables have been widely used to reduce bias from omitted variables in estimates of causal relationships such as the effect of schooling on earnings. Intuitively, instrumental variables methods use only a portion of the variability in key variables to estimate the relationships of interest; if the instruments are valid, that portion is unrelated to the omitted variables. We discuss the mechanics of instrumental variables, and the qualities that make for a good instrument, devoting particular attention to instruments that are derived from "natural experiments." A key feature of the natural experiments approach is the transparency and refutability of identifying assumptions. We also discuss the use of instrumental variables in randomized experiments.

Joshua Angrist
MIT Department of Economics
50 Memorial Drive
Cambridge, MA 02142
and NBER
angrist@mit.edu

Alan B. Krueger
Princeton University
Firestone Library
Princeton, NJ 08544
and NBER
akrueger@princeton.edu

The method of instrumental variables is a signature technique in the econometrics tool kit. The canonical example, and earliest applications, of instrumental variables involved attempts to estimate demand and supply curves.¹ Economists such as P.G. Wright, Henry Schultz, Elmer Working, and Ragnar Frisch were interested in estimating the elasticities of demand and supply for products ranging from herring to butter, usually with time-series data. If both the demand and supply curves shift over time, the observed data on quantities and prices reflect a set of equilibrium points on both curves. Consequently, an ordinary least squares regression of quantities on prices fails to identify – that is, trace out – either the supply or demand relationship.

P.G. Wright (1928) confronted this issue in the seminal application of instrumental variables: estimating the elasticities of supply and demand for flaxseed, the source of linseed oil.² Wright noted the difficulty of obtaining estimates of the elasticities of supply and demand from the relationship between price and quantity alone. He suggested, however, that certain “curve shifters” – what we would now call instrumental variables – can be used to address the problem (p. 312): “Such additional factors may be factors which (A) affect demand conditions without affecting cost conditions or which (B) affect cost conditions without affecting demand conditions.” A variable he used for

1. See Goldberger (1972) and Morgan (1990) for a discussion of the origins of instrumental variables and related methods. Bowden and Turkington (1984) provide a more technical discussion. The term “instrumental variables” originated with Olav Reiersøl (1945); Morgan cites an interview in which Reiersøl attributed the term to his teacher, Ragnar Frisch.

2. In the early 1920s, Wright's son, Sewall Wright, developed “causal path analysis,” a method-of-moments-type technique for estimating recursive structural models and simultaneous equations. P.G. Wright showed that path analysis and instrumental variables were equivalent in

the demand curve shifter was the price of substitute goods, such as cottonseed, while a variable he used for the supply curve shifter was yield per acre, which can be thought of as primarily determined by the weather.

Specifically, an instrumental variables estimate of the demand elasticity can be constructed by dividing the sample covariance between the log quantity of flaxseed and the yield per acre by the sample covariance between the log price of flaxseed and the yield per acre. This estimate is consistent as long as yield per acre is uncorrelated with the error in the demand equation and correlated with price. Replacing yield per acre with the price of substitutes in this calculation generates an instrumental variables estimate of the supply elasticity. Intuitively, weather-related shifts in yield are used to trace out the demand curve, while changes in the price of substitutes are used to shift the demand curve so as to trace out the supply curve.

Wright (1928, p. 314) observed: “Success with this method depends on success in discovering factors of the type A and B.” He used six different supply shifters to estimate the demand curve, and then averaged the six instrumental variables estimates. The resulting average elasticity of demand for flaxseed was $-.80$. His average instrumental variables estimate of the elasticity of supply was 2.4 . Wright’s econometric advance was unnoticed by the subsequent literature. Not until the 1940s were instrumental variables and related methods rediscovered and extended.

Wright’s (1928) method of averaging the different instrumental variables estimates does not necessarily produce the most efficient estimate; other estimators may combine the information in different instruments to produce an estimate with less

his simultaneous equations application. It is quite likely that Sewall Wright deserves much of the credit for his father’s use of instrumental variables.

sampling variability. The most efficient way to combine multiple instruments is usually two-stage least squares, originally developed by Theil (1953).³ In the first stage, the “endogenous” right-hand side variable (price in this application) is regressed on all the instruments. In the second stage, the predicted values of price, based on the data for the instruments and the coefficients estimated from the first-stage regression, are then either plugged directly into the equation of interest in place of the endogenous regressor or, equivalently, used as an instrument for the endogenous regressor. In this way, two-stage least squares takes the information in a set of instruments and neatly boils it down to a single instrument.⁴

Instrumental Variables and Measurement Error

Instrumental variables methods were also pioneered to overcome measurement error problems in explanatory variables.⁵ Measurement error can arise for many reasons, including the limited ability of statistical agencies to collect accurate information, and the deviation between the variables specified in economic theory and those collected in practice. If an explanatory variable is measured with additive random errors, then the

3. The relative efficiency of two-stage least squares turns on a number of auxiliary assumptions, such as homoscedastic errors. See Wooldridge (2001) for a discussion of alternative generalized method of moments estimators.

4. Typically, a number of “exogenous” conditioning variables also appear in both the supply and demand equations. These exogenous covariates do not play the role of instruments but nevertheless should be included in both the first- and second-stage regressions. Two-stage least squares can also be used if there is more than one endogenous regressor in an equation, provided there are at least as many instruments as endogenous regressors (see, for example, Bowden and Turkington, 1984).

5. Wald's (1940) method of fitting straight lines was specifically developed to overcome errors-in-variables problems. Durbin (1954) showed that Wald's method is a special case of instrumental variables. See also Geary (1949). Hausman (2001) provides a recent overview of measurement error problems.

coefficient on that variable in a bivariate ordinary least squares regression will be biased toward zero in a large sample. The higher the proportion of variability that is due to errors, the greater the bias. Given instrument that is uncorrelated with the measurement error and the equation error (that is, the equation error from the model with the correctly measured data) but correlated with the correctly measured variable, then instrumental variables provides a consistent estimate even in the presence of measurement error.

Friedman's (1957) celebrated analysis of the consumption function can be interpreted as an application of instrumental variables in this context. Annual income is a noisy measure of permanent income, so a regression of consumption on annual income yields too small an estimate of the marginal propensity to consume from permanent income. To overcome this measurement problem, Friedman grouped his data by city, which is equivalent to using a two-stage least squares procedure. The first stage is implicitly a regression of annual income on a set of dummies indicating each of the cities. The fitted values from this regression would be average income by city, so that regressing micro consumption data on fitted income values, as is done in two-stage least squares, is the same as a weighted regression using city average data, where the weights are the number of observations per city.

While two-stage least squares and other instrumental variables estimators are consistent, they are not unbiased. Instrumental variables estimates are not unbiased because they involve a ratio of random quantities, for which expectations need not exist or have a simple form. In contrast, expectations of ordinary least squares estimates typically exist and are easily calculated. Textbooks sometimes gloss over the distinction between unbiasedness and consistency, but the difference can matter in practice.

Unbiasedness means the estimator has a sampling distribution centered on the parameter of interest in a sample of any size, while consistency only means that the estimator converges to the population parameter as the sample size grows. Since instrumental variables estimates are consistent, but not unbiased, researchers using instrumental variables should aspire to work with large samples.

The precise form of the asymptotic distribution of an instrumental variables estimator (i.e., the sampling distribution in very large samples) depends on a number of technical conditions. Most modern software packages include options for “robust standard errors” that are asymptotically valid under reasonably general assumptions. It is important to remember, however, that in practice these standard errors are only approximate.

Instrumental Variables and Omitted Variables

Although instrumental variables methods are still widely used to estimate systems of simultaneous equations and to counteract bias from measurement error, a flowering of recent work uses instrumental variables to overcome omitted variables problems in estimates of causal relationships. Studies of this type are primarily concerned with estimating a narrowly defined causal relationship such as the effect of schooling, training, or military service on earnings; the impact of smoking or medical treatments on health; the effect of social insurance programs on labor supply; or the effect of policing on crime. The observed association between the outcome and explanatory variable of interest in these and many other examples is likely to be misleading in the sense that it partly

reflects omitted factors that are related to both variables. If these factors could be measured and held constant in a regression, the omitted variables bias would be eliminated. In practice, however, economic theory typically does not specify all of the variables that should be held constant while estimating a relationship, and it is difficult to accurately measure all of the relevant variables even if they are specified.

One solution to the omitted variable problem is to randomly assign the variable of interest. For example, social experiments are sometimes used to assign people to a job training program or a control group. Random assignment assures that participation in the program (among those in the assignment pool) is not correlated with omitted personal or social factors. Randomized experiments are not always possible, however. It would not be feasible to force a randomly chosen group of people to quit smoking or attend school for an extra year, or to randomly assign the value of the minimum wage across states. On the other hand, it may be possible to find, or even to create, a degree of exogenous variation in variables like schooling, smoking, and minimum wages. Instrumental variables offer a potential solution in these situations.

To see how instrumental variables can solve the omitted variables problem, suppose that we would like to use the following cross-sectional regression equation to measure the “rate of return to schooling,” denoted ρ :

$$Y_i = \alpha + \rho S_i + \beta A_i + \epsilon_i$$

In this equation, Y_i is person i 's log wage, S_i is his or her highest grade of schooling completed, and A_i is a measure of ability or motivation. (For simplicity, we take A_i to be a single variable, although it could be a vector of variables.) Although the problem of estimating this equation is straightforward in principle, data on A_i are typically

unavailable, and researchers are unsure what the right controls for ability or motivation would be in any case.⁶ Without additional information, the parameter of interest, ρ , is not identified; that is, we cannot deduce it from the joint distribution of earnings and schooling alone.

Suppose, however, we have a third variable, the instrument, denoted Z_i , which is correlated with schooling, but otherwise unrelated to earnings. That is, Z_i is uncorrelated with the omitted variables and the regression error, ϵ_i . Then an instrumental variables estimate of the payoff to schooling is the sample analog of $\text{Cov}(Y_i, Z_i)/\text{Cov}(S_i, Z_i)$. The instrumental variables methods allow us to consistently estimate the coefficient of interest free from bias from omitted variables, without actually having data on the omitted variables or even knowing what they are. (If there is more than one valid instrument the coefficient of interest can be estimated by two-stage least squares.) Intuitively, instrumental variables solve the omitted variables problem by using only part of the variability in schooling – specifically, a part that is uncorrelated with the omitted variables – to estimate the relationship between schooling and earnings.

Instruments that are used to overcome omitted variables bias are sometimes said to derive from “natural experiments”.⁷ Recent years have seen a resurgence in the use of instrumental variables in this way; that is, to exploit situations where the forces of nature or government policy have conspired to produce an environment somewhat akin to a randomized experiment. This type of application has generated some of the most

6. The expected coefficient on schooling from a regression that omits the A_i variable is $\rho + \beta\gamma$, where γ is the regression coefficient from a hypothetical regression of A_i on schooling. It should be apparent that if the omitted variable is uncorrelated with schooling, or uncorrelated with earnings conditional on schooling, the coefficient on schooling is an unbiased estimate of ρ even if A_i is omitted from the equation.

provocative empirical findings in economics, along with some controversy over substance and methods.

A good instrument is correlated with the endogenous regressor for reasons the researcher can verify and explain, but uncorrelated with the outcome variable for reasons beyond its effect on the endogenous regressor. Maddala (1977, p. 154) rightfully asks, “Where do you get such a variable?” Like most econometrics texts, he does not provide an answer. In our view, good instruments often come from detailed knowledge of the economic mechanism and institutions determining the regressor of interest.

In the case of schooling, theory suggests that schooling choices are determined by comparing the costs and benefits of alternative choices. Thus, one possible source of instruments would be differences in costs due, say, to loan policies or other subsidies that vary independently of ability or earnings potential. A second source of variation in educational attainment is institutional constraints. Angrist and Krueger (1991) exploit this kind of variation in a paper that typifies the use of “natural experiments” to try to eliminate omitted variables bias.

The rationale for the Angrist and Krueger (1991) approach is that, because most states required students to enter school in the calendar year in which they turned 6, school start age is a function of date of birth. Those born late in the year are young for their grade. In states with a December 31st birthday cutoff, children born in the fourth quarter enter school at age $5 \frac{3}{4}$, while those born in the first quarter enter school at age $6 \frac{3}{4}$. Furthermore, because compulsory schooling laws typically require students to remain in school until their 16th birthday, these groups of students will be in different grades when

7. A natural experiment can be studied without application of instrumental variables methods; in this case, reduced form estimates would be presented.

they reach the legal dropout age. In essence, the combination of school start age policies and compulsory schooling laws creates a natural experiment in which children are compelled to attend school for different lengths of time depending on their birthdays.

Using data from the 1980 census, we looked at the relationship between educational attainment and quarter of birth for men born from 1930 to 1959. The overall pattern is that younger birth cohorts finished more schooling. Figure 1 displays the education-quarter-of-birth pattern for men born in the 1930s. The figure clearly shows that men born early in the calendar year tend to have lower average schooling levels. We selected this 10-year birth cohort because men this age tend to have a relatively flat age-earnings profile. But the pattern of less education for men born early in the year holds for men born in the 1940s and 1950s as well. Because an individual's date of birth is probably unrelated to his or her innate ability, motivation, or family connections (ruling out astrological effects), date of birth should provide a valid instrument for schooling.

Figure 2 displays average earnings by quarter of birth for the same sample. In essence, this figure shows the “reduced form” relationship between the instruments and the dependent variable. Older cohorts tend to have higher earnings, because earnings rise with work experience. But the figure also shows that, on average, men born in early quarters of the year almost always earned less than those born later in the year. Importantly, this reduced form relationship parallels the quarter-of-birth pattern in schooling. An examination of the reduced-form and first-stage estimates, either in graphical or tabular form, often provides insights concerning the causal story motivating a particular set of instrumental variable estimates. In this case, it is clear that the

differences in education and earnings associated with quarter of birth are discrete blips, rather than smooth changes related to the gradual effects of aging.

The intuition behind instrumental variables in this case is that differences in earnings by quarter of birth are assumed to be accounted for solely by differences in schooling by quarter of birth, so that the estimated return to schooling is simply the appropriately rescaled difference in average earnings by quarter of birth. Only a small part of the variability in schooling -- the part associated with quarter of birth -- is used to identify the return to education. In our formal statistical estimates, we found that men born in the first quarter have about one-tenth of a year less schooling than men born in later quarters, and that men born in the first quarter earn about 0.1 percent less than men born in later quarters. The ratio of the difference in earnings to the difference in schooling, about .10, is an instrumental variables estimate of the proportional earnings gain from an additional year of schooling.

As it turns out, this estimate of the change in earnings due to additional education differs little from a simple ordinary least squares regression of education on earnings in our data. This finding suggests that there is little bias from omitted ability variables in the ordinary least squares estimate of the effect of education on earnings, probably because omitted variables in the earnings equation are uncorrelated with education. In other applications, such as Angrist and Lavy's (1999) analysis of the effect of class-size on student achievement based on differences in class-size stemming from Maimonides maximum class size rule, the instrumental variables estimates and ordinary least squares estimates are quite different.

A common criticism of the natural experiments approach to instrumental variables is that it does not fully spell out the underlying theoretical relationships (see, e.g., Rosenzweig and Wolpin, 2000). In the Angrist and Krueger (1991) application, for example, the theoretical relationship between education and earnings is not developed from an elaborate theoretical model; instead, it depends on the institutional details of the education system. Nevertheless, interest in the causal effect of education on earnings is easy to motivate in economics. Moreover, the natural experiments approach to instrumental variables is fundamentally grounded in theory, in the sense that there is usually a well-developed story or model motivating the choice of instruments. Importantly, these stories have implications that can be used to support or refute a behavioral interpretation of the resulting instrumental variables estimates.

For example, the interpretation of the patterns in Figures 1 and 2 as resulting from the interaction of school-start-age policy and compulsory schooling is supported by our finding that quarter of birth is *unrelated* to earnings and educational attainment for those with a college degree or higher. This group is unconstrained by compulsory schooling laws, so if quarter of birth was related to education or earnings in this sample the rational motivating the use of quarter of birth as instruments would have been refuted. The finding that that the identification strategy was not refuted by this test suggests that factors other than compulsory schooling are not responsible for the correlation between education and the instrument in the full sample, and adds credibility to the exercise.

The rich implications and potential refutability of instrumental variables analyses based on natural experiments are an important part of what makes the approach attractive. We would argue that this approach contrasts favorably with studies that provide detailed

but abstract theoretical models, followed by identification based on implausible and unexamined choices about which variables to exclude from the model and assumptions about what statistical distribution certain variables follow. Indeed, one of the most mechanical and naïve, yet common, approaches to the choice of instruments uses atheoretical and hard-to-assess assumptions about dynamic relationships to construct instruments from lagged variables in time series or panel data. The use of lagged endogenous variables as instruments is problematic if the equation error or omitted variables are serially correlated. In this regard, Wright's (1928) use of the more plausible exogenous instrument "yield per acre" seems well ahead of its time.

Interpreting Estimates with Heterogeneous Responses

One difficulty in interpreting instrumental variables estimates is that not every observation's behavior is affected by the instrument. As we have stressed, instrumental variables methods can be thought of as operating by using only part of the variation in an explanatory variable -- that is, by changing the behavior of only some people. For example, in the Angrist and Krueger (1991) study just discussed, the quarter-of-birth instrument is most relevant for those who are at high probability of quitting school as soon as possible, with little or no effect on those who are likely to proceed on to college.

Another example that makes this point is Angrist's (1990) use of Vietnam-era draft lottery numbers as an instrument to estimate the effect of military service on earnings later in life. The draft lottery numbers randomly assigned to young men in the early 1970s were highly correlated with the probability of being drafted into the military,

but not correlated with other factors that might change earnings later. The military draft presumably affected the behavior of those who would not have joined the military otherwise. But most of those who served in Vietnam were true volunteers who would have served regardless of their draft lottery number. An estimate using draft lottery numbers as instruments is therefore based on the experience of draftees only. This may not capture the effects of military service on volunteers' civilian earnings.

In other words, instrumental variables provide an estimate for a specific group -- namely, people whose behavior can be manipulated by the instrument. The quarter-of-birth instruments used by Angrist and Krueger (1991) generate an estimate for those whose level of schooling was changed by that instrument. Similarly, the draft lottery instrument provides estimates of a well-defined causal effect for a subset of the treated group: men whose behavior was changed by the draft lottery "experiment".

This issue arises in many studies using instrumental variables, and it is discussed formally in papers by Imbens and Angrist (1994), and related work. They show that with a dummy endogenous variable, instrumental variables methods estimate causal effects for those whose behavior would be changed by the instrument if it were assigned in a randomized trial. The effect is known as the Local Average Treatment Effect, or LATE for short. In some cases, the experiences of this group of "compliers" are representative of those of the entire treated group. And if everyone in the population has the same response to a particular intervention or treatment, as is commonly assumed, the distinction between LATE and other parameters does not matter. But with

"heterogeneous treatment effects," the parameter identified by instrumental variables may differ from the average effect of interest.⁸

It should be noted that this specificity of estimates is endemic to empirical research. All statistical methods, from the simplest regressions to the most complex structural models, have elements of this limitation when used to analyze phenomena with heterogeneous responses. Nevertheless, many interventions and relationships can be fruitfully studied using estimated effects for specific subsamples, provided the possible limitations to generalizing the results are understood and explored. Indeed, this lack of immediate generality is probably the norm in medical research based on clinical trials, yet much progress has been made in that field.

Our view is that instrumental variables methods often solve the first-order problem of eliminating omitted variables bias for a well-defined population. Since the sample size and range of variability in many empirical studies is quite limited, extrapolation to other populations is naturally somewhat speculative and often relies heavily on theory and common sense. (A fertilizer that helps corn to grow in Iowa will probably have a beneficial effect in California as well, though one can't be sure.) Moreover, the existence of heterogeneous treatment effects would be a reason for analyzing more natural experiments, not fewer, to understand the source and extent of heterogeneity in the effect of interest.

8. The theoretical result that instrumental variables methods identify LATE requires a technical assumption known as "monotonicity". This means that the instrument only moves the endogenous regressor in one direction. With draft lottery instruments, for example, monotonicity implies that being draft-eligible makes a person at least as likely to serve in the military as he would be if he were draft-exempt. This seems reasonable, and is automatically satisfied by traditional latent index models for endogenous treatments.

It is also worth emphasizing that the population one learns about in a natural experiment is often of intrinsic interest. For example, the Angrist and Krueger (1991) instrumental variables estimate, which is identified by differences in schooling for people affected by the compulsory schooling level, is relevant for assessing the economic rewards to increases in schooling induced by legal and institutional changes from policies designed to keep children from dropping out of high school. In the case of the draft lottery, even if the instrumental variables estimates do not necessarily tell us about the effect of military experience on earnings for both volunteers and draftees, knowing the effect of being drafted on later civilian earnings is important. Moreover, the story behind the instrumental variable analysis often opens up other avenues of inquiry. For example, Angrist (1990) interpreted the civilian earnings penalty associated with Vietnam-era service as due to a loss of labor market experience. If true, the resulting estimates have predictive validity for the consequences of compulsory military service in other times and places.

Potential Pitfalls

What can go wrong with instrumental variables? The most important potential problem is a bad instrument, that is, an instrument that is correlated with the omitted variables (or the error term in the structural equation of interest in the case of simultaneous equations). Especially worrisome is the possibility that an association between the instrumental variable and omitted variables can lead to a bias in the resulting estimates that is much greater than the bias in ordinary least squares estimates.

Moreover, seemingly appropriate instruments can turn out to be correlated with omitted variables on closer examination. For example, the weather in Brazil probably shifts the supply curve for coffee, providing a plausible instrument to estimate the effect of price on quantity demanded. But weather in Brazil might also shift the demand for coffee if sophisticated commercial buyers at the New York Coffee, Sugar and Coca Exchange, where coffee futures are traded, use weather data to adjust holdings in anticipation of price increases that may not materialize in fact.

Another concern is the possibility of bias when instruments are only weakly correlated with the endogenous regressor(s). This possibility was first noted by Nagar (1959), and emphasized by Bound, Jaeger and Baker (1995). In fact, instrumental variables estimates with very weak instruments tend to be centered on the corresponding ordinary least squares estimate (Sawa, 1969).

Several solutions to the weak instruments problem have been proposed. First, the bias of two-stage least squares is proportional to the degree of over-identification. In other words, if K instruments are used to estimate the effect of G endogenous variables, the bias is proportional to $K-G$. Using fewer instruments therefore reduces bias. In fact, if the number of instruments is equal to the number of endogenous variables, the bias is approximately zero. A variety of technical fixes and diagnostic tests have also been proposed for the weak instrument problem.⁹

9. One solution is to use the Limited Information Maximum Likelihood (LIML) estimator. Although LIML and two-stage least squares have the same asymptotic distributions and are algebraically equivalent in just-identified models, in over-identified models their finite-sample distributions can be very different. Most importantly, LIML is approximately unbiased in the sense that the median of its sampling distribution is generally close to the population parameter being estimated (Anderson et al., 1982). Alternatives to LIML include the approximately unbiased split-sample and jackknife instrumental variables estimators (Angrist and Krueger, 1995; Angrist, Imbens, and Krueger, 1999; Blomquist and Dahlberg, 1999), bias-corrected

Concerns about weak instruments can be mitigated most simply by looking at the reduced form equation, that is, the ordinary least squares regression of the outcome variable of interest on the instruments and exogenous variables. These estimates are unbiased, even if the instruments are weak. Because the reduced form effects are proportional to the coefficient of interest, one can determine the sign of the coefficient of interest and guesstimate its magnitude by rescaling the reduced form using plausible assumptions about the size of the first-stage coefficient(s). Most importantly, if the reduced form estimates are not significantly different from zero, the presumption should be that the effect of interest is either absent or the instruments are too weak to detect it. The plausibility of the magnitude of the reduced form effects should also be considered.

We conclude our review of pitfalls with a discussion of functional form issues for both the first and second stages in two-stage least squares estimation. For example, researchers are sometimes tempted to use probit or logit for the first stage in two-stage least squares application with a dummy endogenous regressor. But this is not necessary and may even do some harm. In two-stage least squares, consistency of the second stage estimates does not turn on getting the first-stage functional form right (Kelejian, 1971). Moreover, using a nonlinear first stage does not generate consistent estimates unless the nonlinear model happens to be exactly right, so the dangers of mis-specification are high.

Nonlinear second stage estimates with continuous or multi-valued regressors are similarly tricky, requiring a correctly-specified functional form for the estimates to be easily interpreted. And even if the underlying second-stage relationship is nonlinear, linear instrumental variables estimates such as two-stage least squares typically capture

estimation (Sawa, 1973; Bekker, 1994), inference procedures discussed by Staiger and Stock (1997) and Hahn and Hausman (2000), and Bayesian smoothing of the first stage (Chamberlain

an average effect of economic interest analogous to the LATE parameter for dummy endogenous regressors (Angrist, Graddy, and Imbens, 2000; Card, 1995; Heckman and Vytlacil, 2000). Thus, two-stage least squares is a robust estimation method that provides a natural starting point for instrumental variables applications. The importance of functional form issues can be assessed in a more detailed secondary analysis by experimenting with alternative instruments and examining suitable graphs.

Nature's Stream of Experiments

Trygve Haavelmo (1944, p.14) drew an analogy between two sorts of experiments, "those we should like to make", and "the stream of experiments that nature is steadily turning out from her own enormous laboratory, and which we merely watch as passive observers." He also lamented, "unfortunately -- most economists do not describe their designs of experiments explicitly." The defining characteristic of many recent applications of instrumental variables to the omitted variables problem is the attention devoted to describing and assessing the underlying quasi-experimental design. This can be seen as an explicit attempt to use observational data to mimic randomized experiments as closely as possible.

Some economists are pessimistic about the prospects of finding a substantial number of useful natural experiments. Michael Hurd (1990), for example, called the search for natural experiments to test the effect of Social Security on labor supply "overly cautious" and warned, "if applied to other areas of empirical work [this method] would effectively stop estimation." We make no claim that natural experiments are the only way

and Imbens, 1996).

to obtain useful results, only that they have the potential to greatly increase our understanding of important economic relationships. Table 1 provides a sampling of some recent studies that have used instrumental variables to analyze a natural experiment or a researcher-generated randomized experiment. It is hard to conclude that empirical work has effectively stopped.

The first panel in Table 1 illustrates the breadth of application of the natural experiments idea in recent empirical work. Other examples can be found in the surveys by Meyer (1995) and Rosenzweig and Wolpin (2000). Some of the examples are more convincing than others. But all are distinguished by a serious attempt to substantiate the underlying assumptions used to infer causality. There is more "theory" behind these attempts than in many ostensibly structural models where the justification for including or excluding certain variables is neither explicitly described nor evaluated.

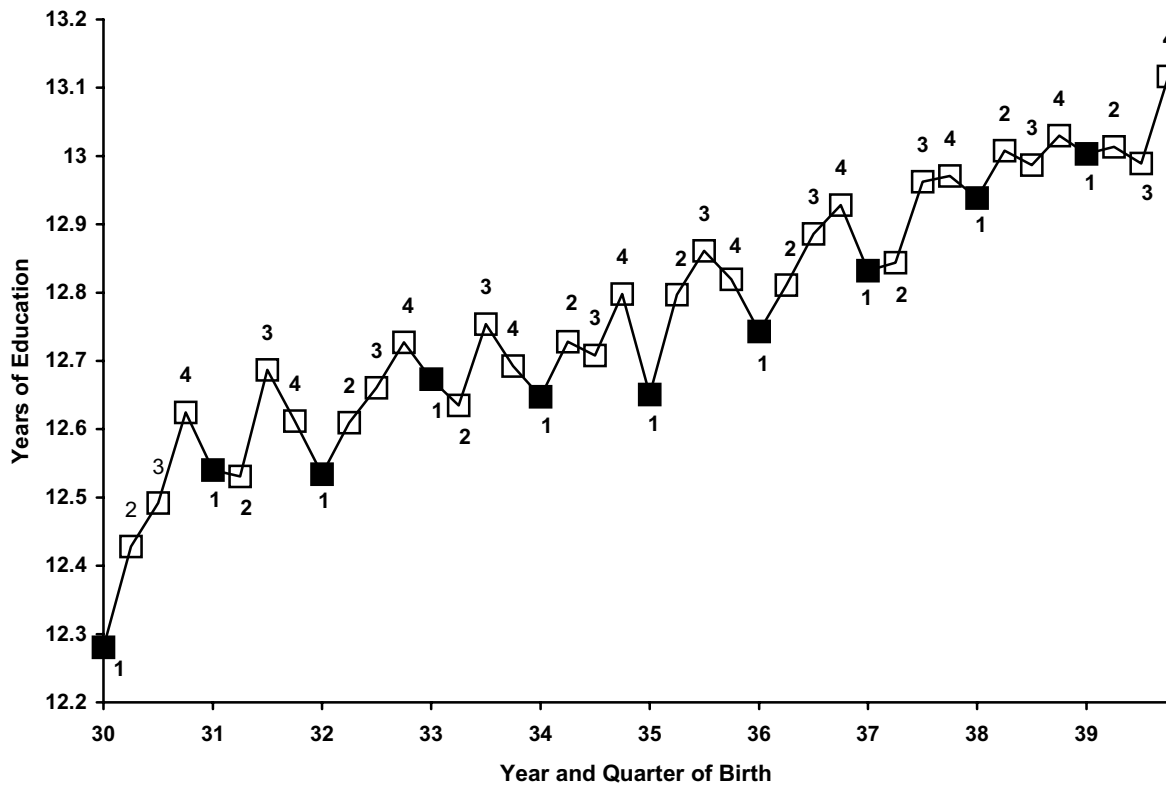
The second panel in Table 1 illustrates another important development: the use of instrumental variables in randomized experiments. Instrumental variables are useful in experiments when, either because of practical or ethical considerations, there is incomplete compliance in the treatment or control groups. In randomized evaluations of training programs, for example, some treatment group members may decline training while some control group members may avail themselves of training through channels outside the experiment. Similarly, in medical trials, doctors may be willing to randomly offer, but not impose, incentives that change behaviors like smoking or taking a new medication.

Even in experiments with compliance problems, instrumental variables can be used to estimate the effect of interventions such as job training or medical treatments.

The instrumental variable in such cases is a dummy variable indicating randomized assignment to the treatment or control group and the endogenous right-hand-side variable is an indicator of actual treatment status. For example, the actual treatment status variable would be a dummy variable that equals one for each treatment and control group member who participated in training, and zero for all those who did not participate in training. This approach yields a consistent estimate of the causal effect of the treatment for the population that complies with their random assignment, i.e., the population of “compliers” (see Imbens and Angrist, 1994). As in natural experiments, the instrument is used to exploit an exogenous source of variation – created by explicit random assignment in these cases – to estimate the effect of interest. The use of such researcher-generated instruments is growing and reflects the accelerating convergence of classical experimentation and observational research methods.


Our view is that progress in the application of instrumental variables methods depends mostly on the gritty work of finding or creating plausible experiments that can be used to measure important economic relationships -- what statistician David Freedman (1991) has called "shoe-leather" research. Here the challenges are not primarily technical in the sense of requiring new theorems or estimators. Rather, progress comes from detailed institutional knowledge, and the careful investigation and quantification of the forces at work. Of course, such endeavors are not really new. They have always been at the heart of good empirical research.

Figure 1: Mean Years of Completed Schooling by Year and Quarter of Birth



Source: Authors' calculations based on 1980 Census

Figure 2: Mean Log Weekly Earnings, by Year and Quarter of Birth



Source: Authors' calculations from the 1980 Census

Table 1: Examples of studies that use instrumental variables to analyze data from natural and randomized experiments

Outcome Variable	Endogenous Variable	Source of Instrumental Variable(s)	Reference
1. Natural Experiments			
Labor supply	Disability insurance Replacement rates	Region and time variation in benefit rules	Gruber (2000)
Labor supply	Fertility	Sibling-Sex composition	Angrist and Evans (1998)
Education, Labor Supply	Out-of-wedlock fertility	Occurrence of twin births	Bronars and Grogger (1994)
Wages	UI tax rate	State laws	Anderson and Meyer (2000)
Earnings	Years of schooling	Region and time variation In school construction	Duflo (2001)
Earnings	Years of schooling	Proximity to college	Card (1995)
Earnings	Years of schooling	Quarter of birth	Angrist and Krueger (1991)
Earnings	Veteran status	Cohort dummies	Imbens and van der Klaauw (1995)
Earnings	Veteran status	Draft lottery number	Angrist (1990)
Achievement test scores	Class size	Discontinuities in class size due to maximum class-size rule	Angrist and Lavy (1999)
College enrollment	Financial aid	Discontinuities in financial aid formula	van der Klaauw (1996)
Health	Heart attack Surgery	Proximity to cardiac care centers	McClellan, <i>et al</i> (1994)
Crime	Police	Electoral cycles	Levitt (1997)

Table 1 (Continued)

Outcome Variable	Endogenous Variable	Source of Instrumental Variable(s)	Reference
Employment and Earnings	Length of prison sentence	Randomly assigned Federal judges	Kling (1999)
Birth weight	Maternal Smoking	State cigarette taxes	Evans and Ringel (1999)
2. Randomized Experiments			
Earnings	Participation in federal Job training program	Random assignment of admission to training program	Bloom, et al. (1997)
Earnings	Participation in federal Job Corps program	Random assignment of admission to training program	Schochet, et al. (1997)
Achievement test scores	Enrollment in private School	Randomly selected offer of school voucher	Howell, et al. (2000)
Achievement test scores	Class size	Random assignment to a small or normal-size class	Krueger (1999)
Achievement test scores	hours of study	Random mailing of test preparation materials	Powers and Swinton (1984)
Birth weight	Maternal Smoking	Random assignment of free smoker's counseling	Permutt and Hebel (1989)

REFERENCES

- Anderson, Patricia, and Bruce Meyer. 2000. "The Effects of the Unemployment Insurance Payroll Tax on Wages, Employment, Claims, and Denials," *Journal of Public Economics* 78, 81-106.
- Anderson, T. W., Naoto Kunitomo and Takamitsu Sawa. 1982. "Evaluation of the Distribution Function of the Limited Information Maximum Likelihood Estimator," *Econometrica* 50, 1009-1027.
- Angrist, Joshua D. 1990. "Lifetime Earnings and the Vietnam Era Draft Lottery: Evidence from Social Security Administrative Records," *American Economic Review* 80, 313-335.
- Angrist, Joshua D. and William N. Evans. 1998. "Children and their Parents' Labor Supply: Evidence from Exogenous Variation in Family Size," *American Economic Review* 88, 450-477.
- Angrist, Joshua D. Katy Graddy, and Guido W. Imbens. 2000. "Nonparametric Demand Analysis with an Application to the Demand for Fish," *Review of Economic Studies* 67(3), 499-527.
- Angrist, Joshua D., Guido W. Imbens and Alan B. Krueger. 1999. "Jackknife Instrumental Variables Estimation," *Journal of Applied Econometrics* 14(1), 57-67.
- Angrist, Joshua D. and Alan B. Krueger. 1991. "Does Compulsory School Attendance Affect Schooling and Earnings?" *Quarterly Journal of Economics* 106, 979-1014.
- Angrist, Joshua D. and Alan B. Krueger. 1995. "Split-Sample Instrumental Variables Estimates of the Returns to Schooling," *Journal of Business and Economic Statistics* 13(2), 225-35.
- Angrist, Joshua D. and Alan B. Krueger. 1999. "Empirical Strategies in Labor Economics," *The Handbook of Labor Economics*, Volume III, Amsterdam: North-Holland.
- Angrist, Joshua D. and Victor Lavy. 1999. "Using Maimonides' Rule to Estimate the Effects of Class Size on Scholastic Achievement," *Quarterly Journal of Economics* 104, 533-576.
- Bekker, Paul A. 1994. "Alternative Approximations to the Distributions of Instrumental Variables Estimators," *Econometrica* 62(3), 657-81.
- Blank, Rebecca. 1991. "The Effects of Double-Blind versus Single-blind Reviewing: Experimental Evidence from the American Economic Review." *American Economic Review* 81(5), 1041-67.

- Blomquist, Sören and Matz Dahlberg. 1999. "Small Sample Properties of LIML and Jackknife IV Estimators: Experiments with Weak Instruments," *Journal of Applied Econometrics* 14, 69-88.
- Bloom, H.S., L.L. Orr, S.H. Bell, G. Cave, F. Doolittle, W. Lin and J.M. Bos. 1997. "The Benefits and Costs of JTPA Title II-A Programs," *The Journal of Human Resources* 32, 549-576.
- Bound, John, David Jaeger and Regina Baker. 1995. "Problems with Instrumental Variables Estimation when the Correlation Between the Instruments and the Endogenous Explanatory Variable is Weak," *Journal of the American Statistical Association* 90(430), 443-50.
- Bowden, Roger J. and Darrell A. Turkington. 1984. *Instrumental Variables*, Cambridge: Cambridge University Press.
- Card, David E. 1995. "Earnings, Schooling and Ability Revisited," in Solomon W. Polachek, ed. *Research in Labor Economics* 14, 23-48.
- Chamberlain, Gary and Guido W. Imbens. 1996. "Hierarchical Bayes Models with Many Instrumental Variables," Harvard University Department of Economics, Discussion Paper No. 1781, December.
- Duflo, Esther. 2001. "Schooling and Labor Market Consequences of School Construction in Indonesia: Evidence from an Unusual Policy Experiment," *American Economic Review*, forthcoming.
- Durbin, J. 1954. "Errors in Variables," *Review of the International Statistical Institute* 22, 23-32.
- Evans, William N. and Jeanne Ringel. 1999. "Can Higher Cigarette Taxes Improve Birth Outcomes?," *Journal of Public Economics* 72,135-54.
- Freedman, David. 1991. "Statistical Models and Shoe Leather," in Peter Marsden, ed., *Sociological Methodology 1991*. Chapter 10. American Sociological Association: Washington, DC.
- Friedman, Milton. 1957. *A Theory of the Consumption Function*. Princeton: Princeton University Press.
- Geary, R.C. 1949. "Determination of Linear relations Between Systematic Parts of Variables with Errors of Observations, the Variances of Which are Unknown," *Econometrica* 46, 225-8.
- Goldberger, Arthur S. 1972. "Structural Equation Methods in the Social Sciences," *Econometrica* 40, 979-1001.

- Gruber, Jonathan. 2000. "Disability Insurance Benefits and Labor Supply," *Journal of Political Economy* 108, 1162-1183.
- Haavelmo, T. 1944. "The Probability Approach in Econometrics," *Econometrica* 12 (Supplement), 1-115.
- Hahn, Jinyong, and Jerry Hausman. 1999. "A New Specification Test for the Validity of Instrumental Variables," MIT Working Paper No. 99-10, May.
- Hausman, Jerry. 2001. "Mismeasured Variables in Econometric Analysis: Problems from the Right and Problems from the Left," *Journal of Economic Perspectives*, forthcoming.
- Heckman, James, and Edward Vytlacil. 1999. "Instrumental Variables Methods for the Correlated Random Coefficient Model: Estimating the Average Rate of Return to Schooling when the Return is Correlated with Schooling," *Journal of Human Resources* 33(4), 974-987.
- Howell, W.G., P.J. Wolf, P. E. Peterson, and D.E, Campbell 2000. "Test-Score Effects of School Vouchers in Dayton, New York, and Washington: Evidence from Randomized Field Trials," paper presented at the annual meeting of the American Political Science Association, Washington, D.C., September.
- Hurd, Michael D. 1990. "Research on the Elderly: Economic Status, Retirement, and Consumption and Savings," *Journal of Economic Literature* 28(2), 565-637.
- Imbens, Guido W. and Joshua D. Angrist. 1994. "Identification and Estimation of Local Average Treatment Effects," *Econometrica* 62(2), 467-75.
- Kelejian, H.H. 1971. "Two-Stage Least Squares and Econometric Systems Linear in Parameters but Nonlinear in the Endogenous Variables," *Journal of the American Statistical Association* 66, 373-374.
- Kling, Jeffrey R. 1999. "The Effect of Prison Sentence Length on the Subsequent Employment and Earnings of Criminal Defendants," Princeton University Woodrow Wilson School, Discussion Paper 208.
- Krueger, Alan B. 1999. "Experimental Estimates of Education Production Functions," *Quarterly Journal of Economics* 104, 497-532.
- Levitt, Steven D. 1997. "Using Electoral Cycles in Police Hiring to Estimate the Effect of Police on Crime," *American Economic Review* 87, 270-290.
- Maddala, G.S. 1977. *Econometrics*, New York: McGraw-Hill, Inc.

McClellan, Mark, Barbara J. McNeil, and Joseph P. Newhouse. 1994. "Does More Intensive Treatment of Acute Myocardial Infarction in the Elderly Reduce Mortality?," *Journal of the American Medical Association* 272, 859-893.

Meyer, Bruce D. 1995. "Natural and Quasi-experiments in Economics," *Journal of Business and Economic Statistics* 13(2), 151-61.

Meyer, Bruce D. W. Kip Viscusi and David L. Durbin. 1995. "Workers' Compensation and Injury Duration: Evidence from a Natural Experiment," *American Economic Review* 85, 322-40.

Morgan, Mary S. 1990. *The History of Econometric Ideas*, Cambridge: Cambridge University Press.

Nagar, A.L. (1959), "The Bias and Moment Matrix of the General k-Class Estimators of the Parameters in Simultaneous Equations," *Econometrica* 27, 575-595.

Permutt, Thomas, and J. Richard Hebel. 1989. "Simultaneous Equation Estimation in a Clinical Trial of the Effect of Smoking on Birth Weight," *Biometrics* 45, 619-22.

Powers, Donald E., and Spencer S. Swinton. 1984. "Effects of Self-Study for Coachable Test Item Types," *Journal of Educational Psychology* 76, 266-278.

Reiersøl, Olav. 1945. "Confluence Analysis by Means of Instrumental Sets of Variables," *Arkiv för Matematik, Astronomi och Fysik*, 32a(4), 1-119.

Rosenzweig, Mark R. and Kenneth I. Wolpin. 2000. "Natural 'Natural Experiments' in Economics," *Journal of Economic Literature* 38, 827-874.

Sawa, Takamitsu. 1969. "The Exact Sampling Distribution of Ordinary Least Squares and Two-Stage Least Squares Estimators," *Journal of the American Statistical Association* 64(327), 923-937.

Sawa, Takamitsu. 1973. "An Almost Unbiased Estimator in Simultaneous Equations Systems," *International Economic Review* 14(1), 97-106.

Schochet, Peter, John Burghardt, and Steven Glazerman. 1999. "The National Job Corps Study: Short Term Impacts of Job Corps on Participants' Employment and Related Outcomes," Mathematica Corp., Princeton, New Jersey.

Staiger, Douglas and James H. Stock. 1997. "Instrumental Variables Regression with Weak Instruments," *Econometrica* 65(3), 557-86.

Theil, H. 1953. "Repeated Least Squares Applied to Complete Equation Systems," The Hague: Central Planning Bureau.

Van der Klaauw. 1996. "A Regression-Discontinuity Evaluation of the Effect of Financial Aid Offers on College Enrollment," Department of Economics, New York University, mimeo.

Wooldridge, Jeffrey. "Applications of Generalized Method of Moments Estimation," this volume.

Wright, Phillip G. 1928. *The Tariff on Animal and Vegetable Oils*," New York: Macmillan.