Journal of Econometrics 156 (2010) 3–20

Contents lists available at ScienceDirect

Journal of Econometrics

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

Structural vs. atheoretic approaches to econometrics

Michael P. Keane *

University of Technology Sydney, Australia
Arizona State University, United States

ARTICLE INFO

Article history:
Available online 16 September 2009

JEL classification:
B23
C10
C21
C52
J24

Keywords:
Structural models
Natural experiments
Dynamic models
Life-cycle models
Instrumental variables

ABSTRACT

In this paper I attempt to lay out the sources of conflict between the so-called "structural" and "experimentalist" camps in econometrics. Critics of the structural approach often assert that it produces results that rely on too many assumptions to be credible, and that the experimentalist approach provides an alternative that relies on fewer assumptions. Here, I argue that this is a false dichotomy. All econometric work relies heavily on *a priori* assumptions. The main difference between structural and experimental (or "atheoretic") approaches is not in the number of assumptions but the extent to which they are made explicit.

1. Introduction

The goal of this volume is to draw attention to the many researchers, especially young researchers, doing high quality structural econometric work in several areas of applied microeconomics. It is motivated by a perception that structural work has fallen out of favor in recent years, and that, as a result, the work being done by such young researchers has received too little attention. In this paper, I would like to talk about why structural work has fallen out of favor, whether that ought to be the case, and, if not, what can be done about it. I will argue that there is much room for optimism, as recent structural work has increased our understanding of many key issues.

Since roughly the early 90s, a so-called "experimentalist" approach to econometrics has been in vogue. This approach is well described by Angrist and Krueger (1999), who write that "Research in a structuralist style relies heavily on economic theory to guide empirical work ... An alternative to structural modeling, ... the "experimentalist" approach, ... puts front and center the problem of identifying causal effects from specific events or situations". By "events or situations", they are referring to "natural experiments" that generate exogenous variation in certain variables that would otherwise be endogenous in the behavioral relationship of interest.

The basic idea here is this. Suppose we are interested in the effect of a variable X on an outcome Y, for example, the effect of an additional year of education on earnings. The view of the "experimentalist" school is that this question is very difficult to address precisely because education is not randomly assigned. People with different education levels tend to have different levels of other variables U, at least some of which are unobserved (e.g., innate ability), that also affect earnings. Thus, the "causal effect" of an additional year of education is hard to isolate.

However, the experimentalist school seems to offer us a way out of this difficult problem. If we can find an "instrumental variable" Z that is correlated with X but uncorrelated with the unobservables that also affect earnings, then we can use an instrumental variable (IV) procedure to estimate the effect of X on Y. The "ideal instrument" is a "natural experiment" that generates random assignment (or something that resembles it), whereby those with $Z = 1$ tend, *Ceteris paribus*, to chose a higher level of X than those with $Z = 0$. That is, some naturally occurring event affects a random subset of the population, inducing at least some members of that "treatment group" to choose or be assigned a higher level of X than they would have otherwise.[1] *Prima facie*, this approach

* Corresponding address: University of Technology Sydney, PO Box 123, Broadway, NSW 2007, Australia. Tel.: +61 2 9514 9742; fax: +61 2 9514 9743.
E-mail address: michael.keane@uts.edu.au.

[1] As Angrist and Krueger (1999) state: "In labor economics at least, the current popularity of quasi-experiments stems ... from this concern: Because it is typically impossible to adequately control for all relevant variables, it is often desirable to seek situations where it is reasonable to presume that the omitted variables are uncorrelated with the variables of interest. Such situations may arise if ... the forces of nature or human institutions provide something close to random assignment".

0304-4076/$ – see front matter © 2009 Elsevier B.V. All rights reserved.
doi:10.1016/j.jeconom.2009.09.003

does not seem to require strong assumptions about how economic agents chose X, or how U is generated.

This seemingly simple idea has found widespread appeal in the economics profession. It has led to the currently prevalent view that, if we can just find "natural experiments" or "clever instruments", we can learn interesting things about behavior without making strong *a priori* assumptions, and without using "too much" economic theory. In fact, I have heard it said that: "empirical work is all about finding good instruments", and that, conversely, results of structural econometric analysis cannot be trusted because they hinge on "too many assumptions". These notions seem to account for both the current popularity of atheoretic approaches to econometrics, and the relative disfavor into which structural work has fallen.

Here, I want to challenge the popular view that "natural experiments" offer a simple, robust and relatively "assumption free" way to learn interesting things about economic relationships. Indeed, I will argue that it is not possible to learn anything of interest from data without theoretical assumptions, even when one has available an "ideal instrument".[2] Data cannot determine interesting economic relationships without *a priori* identifying assumptions, regardless of what sort of idealized experiments, "natural experiments" or "quasi-experiments" are present in that data.[3] Economic models are always needed to provide a window through which we interpret data, and our interpretation will always be subjective, in the sense that it is contingent on our model.

Furthermore, atheoretical "experimentalist" approaches do not rely on fewer or weaker assumptions than do structural approaches. The real distinction is that, in a structural approach, one's *a priori* assumptions about behavior must be laid out explicitly, while in an experimentalist approach, key assumptions are left implicit. I will provide some examples of the strong implicit assumptions that underlie certain "simple" estimators to illustrate this point.

Of course, this point is not new. For instance, Heckman (1997) and Rosenzweig and Wolpin (2000) provide excellent discussions of the strong implicit assumptions that underlie conclusions from experimentalist studies, accompanied by many useful examples. Nevertheless, the perception that experimental approaches allow us to draw inferences without "too much" theory seems to stubbornly persist. Thus, it seems worthwhile to continue to stress the fallacy of this view. One thing I will try to do differently from the earlier critiques is to present even simpler examples. Some of these examples are new, and I hope they will be persuasive to a target audience that does not yet have much formal training in either structural or experimentalist econometric approaches (e.g., first year graduate students).

If one accepts that inferences drawn from experimentalist work are just as contingent on *a priori* assumptions as those from structural work, the key presumed advantage of the experimentalist approach disappears. One is forced to accept that all empirical work in economics, whether "experimentalist" or "structural", relies critically on *a priori* theoretical assumptions. But once we accept the key role of *a priori* assumptions and the inevitability of subjectivity in all inference, how can we make more progress in applied work in general?

I will argue that this key role of *a priori* theory in empirical work is not really a problem – its something economics has in common with other sciences – and that, once we recognize the contingency of all inference, it becomes apparent that structural, experimentalist and descriptive empirical work all have complimentary roles to play in advancing economics as a science. Finally, I will turn to a critique of prior work in the structural genre itself. I will argue that structural econometricians need to devote much more effort to validating structural models, a point previously stressed in Wolpin (1996) and Keane and Wolpin (1997a,b, 2007). This is a difficult area, but I will describe how I think progress can be made.

2. Even "ideal" instruments tell us nothing without *a priori* assumptions

When I argue we cannot ever learn anything from natural experiments without *a priori* theoretical assumptions, a response I often get, even from structural econometricians, is this: "you have to concede that when you have an ideal instrument, like a lottery number, results based on it are incontrovertible". In fact, this is a serious misconception that needs to be refuted. One of the key papers that marked the rising popularity of the experimentalist approach was Angrist (1990), who used Vietnam era draft lottery numbers – which were randomly assigned but influenced the probability of "treatment" (i.e., military service) – as an instrument to estimate the effect of military service on subsequent earnings. This paper provides an excellent illustration of just how little can be learned without theory, even when we have such an "ideal" instrument.

A simple description of that paper is as follows: The sample consisted of men born from '50–'53. The 1970 lottery affected men born in '50; the '71 lottery affected men born in '51, etc. Each man was assigned a lottery number from 1 to 365 based on random drawings of birth dates, and only those with numbers below a certain ceiling (e.g., 95 in 1972) were draft eligible. Various tests and physical exams were then used to determine the subset of draft eligible men who were actually drafted into the military (which turned out to be about 15%). Thus, for each cohort, Angrist runs a regression of earnings in some subsequent year ('81 through '84) on a constant and a dummy variable for veteran status. The instruments are a constant and a dummy variable for draft eligibility. Since there are two groups, this leads to the Wald (1940) estimator, $\hat{\beta} = (\bar{y}^E - \bar{y}^N)/(P^E - P^N)$, where $\bar{y}^E$ denotes average earnings among the draft eligible group, P^E denotes the probability of military service for members of the eligible group, and $\bar{y}^N$ and P^N are the corresponding values for the non-eligible group. The estimates imply that military service reduced annual earnings for whites by about \$1500 to \$3000 in 1978 dollars (with no effect for blacks), about a 15% decrease. The conclusion is that military service actually lowered earnings (i.e., veterans did not simply have lower earnings because they tended to have lower values of the error term U to begin with).

[2] By "data" I mean the joint distribution of observed variables. To use the language of the Cowles Commission, "Suppose . . . B is faced with the problem of identifying . . . the structural equations that alone reflect specified laws of economic behavior . . . Statistical observation will in favorable circumstances permit him to estimate . . . the probability distribution of the variables. Under no circumstances whatever will passive statistical observation permit him to distinguish between different mathematically equivalent ways of writing down that distribution . . . . The only way in which he can hope to identify and measure individual structural equations . . . is with the help of a priori specifications of the form of each structural equation" — see Koopmans et al. (1950).

[3] The term "quasi-experiment" was developed in the classic work by Campbell and Stanley (1963). In the quasi-experiment, unlike a true experiment, subjects are not randomly assigned to treatment and control groups by the investigator. Rather, events that occur naturally in the field, such as administrative/legislative fiat, assign subjects to treatment and control groups. The ideal is that these groups appear very similar prior to the intervention, so that the event in the field closely resembles randomization. To gauge pre-treatment similarity, it is obviously necessary that the data contain a pre-treatment measure for the outcome of interest. Campbell and Stanley (1963) list several other types of research designs based on observational data which do not satisfy this criterion, such as studies based on "one-shot" cross-section surveys, which do not provide a pre-treatment outcome measure. They also emphasize that, even when treatment and control groups are very similar on observables prior to treatment, they may differ greatly on unobservables, making causal inferences from a quasi-experiment less clear than those from a true experiment.