Introduction to the JBES Special Issue on Structural Estimation in Applied Microeconomics

Keane, Michael and Wolpin, Kenneth

1997

Introduction to the *JBES* Special Issue on Structural Estimation in Applied Microeconomics

Michael P. KEANE
Department of Economics, University of Minnesota, Minneapolis, MN 55455 (mkeane@atlas.socsci.umn.edu)

Kenneth I. WOLPIN
Department of Economics, University of Pennsylvania, Philadelphia, PA 19104

1. INTRODUCTION

All of the articles in this volume use economic theory to guide empirical work in applied microeconomics. They span several empirical methodologies that differ, most importantly, by whether the parameters that are estimated are interpretable as the primitives of economic theory (i.e., preferences and technology). Many of the articles have as their goal the recovery of these primitives and are thus explicitly exercises in structural estimation. Others seek to estimate decision rules derived from economic models. But all of the articles illustrate the important role of economic theory in drawing inferences about behavior from data.

Several approaches to structural estimation are represented in this volume. Some of the authors use a *full-solution* approach based on complete solution of the optimization problem confronting agents (Bernard Elyakime, Jean-Jacques Laffont, Patrice Loise, and Quang Vuong; Marcel Fafchamps and John Pender; Christopher Ferrall; Christopher J. Flinn; Tony Lancaster; Geoffrey Rothwell and John Rust). Others identify primitives from the first-order conditions of the optimization problem (Richard Blundell, Thierry Magnac, and Costas Meghir; Robert A. Miller and Holger Seig) or from exact reduced-form representations of the decision rules of the economic agents (Gerard J. van den Berg and Cees Gorter). A variety of reduced-form approaches are also represented. One approach is to estimate the reduced-form parameters of the exact decision rules (Lee Lillard and Yoram Weiss). An alternative is to use the theory to specify approximations to the reduced-form decision rules and then estimate the parameters of these approximations (Jere R. Behrman, Andrew Foster, and Mark R. Rosenzweig; R. Mark Gritz and Thomas MaCurdy).

It is useful to contrast these various approaches to estimation. The full-solution approach to structural estimation requires that the optimization problem be solved (usually numerically) for the optimal decision rules of the agents in the model. Implementation of this empirical methodology entails a search over the set of feasible values for the primitives such that the resultant optimal decision rules generate behavior that resembles observed data. The full-solution methodology is in principle the most general approach to structural estimation. The cost of this generality is computational burden (i.e., it requires repeated solution of the optimization problem at each trial parameter value) and the resulting constraints on the complexity of the economic models that are feasible to estimate. An important benefit is that one can forecast the behavior of agents given any change in the state of the world that can be characterized as a change in their constraints—including changes along dimensions that are invariant in the data (see Wolpin 1996).

In some settings, structural estimation can be based on the first-order conditions (FOC's) of the optimization problem. Because FOC's can be much simpler to derive than optimal decision rules and can often be represented analytically, the estimation procedure is generally less computationally demanding than the full-solution method. The gain in computational ease must be balanced against some important drawbacks, however. Clearly, lack of continuity in the agents' choice variables (e.g., discrete choices) raises problems for estimation based on FOC's. Moreover, in most cases not all of the primitive parameters can be identified from the FOC's. Policy experiments often require knowledge of all the primitives.

In some cases analytical representations of the reduced-form decision rules can be obtained with suitable assumptions. The advantage of estimating reduced-form decision rules is that it does not require repeated solution of the economic model. A disadvantage is that in only rare cases are the primitive (structural) parameters exactly identified from the reduced-form parameters. The absence of structural parameter estimates precludes using the results to forecast the outcomes of policy experiments that involve changes in state variables that were invariant in the data used for estimation.

In most economic models the decision rules are highly nonlinear functions of the primitive parameters and of the variables that characterize the state of the world—the *state variables*—including the stochastic elements unobserved by the researcher. Thus, a truly reduced-form approach to estimation of the decision rules is often intractable. An alternative to estimation of the exact reduced-form decision rules is to estimate approximate decision rules. In this approach, the theoretical model determines which state variables enter the optimal decision rule for each choice variable, and decision rules are approximated by some flexible form (e.g., as polynomials in the relevant state variables). But, as with estimation of exact decision rules, an important limitation is that policy experiments can only be performed if they can

Journal of Business & Economic Statistics
April 1997, Vol. 15, No. 2

be characterized as changes in state variables that varied in the estimation data.

All of these estimation approaches, exemplified by the articles in this volume, rely on economic models to interpret the parameter estimates. Which method of estimation is adopted to address a particular substantive issue will depend on an assessment of their relative merits. As these articles demonstrate, different researchers come to different judgments even for the same substantive issue.

The theory-guided approach to estimation represented in this volume is far from being universally accepted in empirical economic research. In fact, it has recently become quite fashionable to eschew theory in favor of a "data should speak for themselves" view. According to this perspective, empirical researchers should rely on randomized social experiments or hunt for "natural experiments" from which inferences about "treatment effects" can supposedly be drawn without reliance on economic theory and without having to make strong auxiliary assumptions [e.g., see the recent *JBES* Symposium on Program and Policy Evaluation (Angrist 1995)]. But, although experiments can tell us what happened to a particular population at a particular time for a specific treatment, to generalize to other contexts requires that we learn about the mechanism generating the response. Thus, generalizing the results of randomized social and/or "natural" experiments beyond the specific context in which they occurred would first require that behavioral theories be posited that are consistent with the (quasi) experimental results and second would require that further evidence be gathered to cast light on the ability of those theories to predict behavior in other contexts. Articles that adopted such an approach would fit the theme of this volume.

2. THE ARTICLES

The articles in this volume serve to illustrate convincingly its theme, that there is a valuable synergistic relationship between economic theory and empirical work in economics. There are several ways to group the articles—by field, by estimation approach, by structural aspects of the theoretical model, and so forth. The articles span industrial organization, labor economics, economic development, and public economics; the estimation approaches are structural and nonstructural; and the theoretical frameworks encompass single-agent and multiagent optimization problems. None of these classifications by itself, however, would serve to demonstrate the contributions of the articles to the overall theme of the volume, although they could serve to highlight their contributions more generally. We therefore have taken the liberty to group the articles in ways that cross these boundaries.

Two of the articles estimate models of job-search behavior, adopting structural and reduced-form approaches, respectively. Ferrall structurally estimates models of labor-force transitions on Canadian and U.S. data. He uses the models to assess the effect of the unemployment insurance (UI) system in Canada on transitions from school to work. A key feature of the Canadian UI system is that new labor-market entrants must hold a first job for 20 weeks before they are eligible for UI. Ferrall's results indicate that unemployment durations for new labor-market entrants in Canada would be quite sensitive to the length of this eligibility requirement. Note that UI eligibility requirements are uniform throughout Canada, thus precluding the use of provincial variation to identify the effect of lowering the eligibility requirement.

In contrast, Gritz and MaCurdy assess the impact of the U.S. UI system on nonemployment durations of youths by estimating approximate decision rules and exploiting interstate variation in UI parameters for identification. They consider the effects of UI on the length of nonemployment spells and the classification of these spells into unemployment and out-of-the-labor-force and the effect of the generosity of UI on the "take-up" rate. Their empirical results imply that weeks of eligibility is a much stronger determinant of nonemployment-spell characteristics than is the level of benefits.

There are also two methodological articles that address issues of parameter identification in models of labor-force transitions. Blundell, Magnac, and Meghir rigorously examine the identification of the primitives in a model of intertemporal consumption and labor-force transitions. Models that link unemployment due to labor-market frictions with consumption smoothing and precautionary savings motives for asset accumulation have not been studied in detail. They show that, with data on assets, it is possible under certain conditions to identify job-offer arrival and job layoff rates as well as the structure of preferences without resorting to a full solution of the agents' dynamic optimization problem.

Lancaster considers Bayesian inference about the primitive parameters in the standard infinite-horizon job-search model. In one specification he treats the reservation wage as a latent variable and does not impose the restriction that it is derived from solution of the wealth-maximization problem. In another specification, the restriction implied by wealth maximization is imposed. In a set of simulation experiments, he finds that, if the observed data are only unemployment durations and accepted wages (the usual case), then imposing the restriction enables one to identify the discount rate, but it conveys little additional information about the job-offer arrival rate and the parameters of the wage-offer distribution, which are identified without it. The implication is that the wage and duration data alone are quite informative about these parameters (given the assumption that a reservation wage rule is being followed).

The articles by Fafchamps and Pender and by Rothwell and Rust illustrate the breadth of applications to which full-solution structural estimation can be applied in single-agent models. Fafchamps and Pender estimate a model of precautionary savings in the context of Indian agriculture, in which there are important nondivisible and irreversible investment opportunities. They find that nondivisibility alone inhibits poor farmers from self-financing what would otherwise appear to be profitable investments; at their income levels irreversibility has little further deterrent effect.

Rothwell and Rust estimate a model of optimal operating decisions for nuclear power plants with operations at vari-

ous levels of capacity, shutdowns for refueling or maintenance, and decommissioning of the plant as options. They use the model to predict the impact of a policy intervention in which the Nuclear Regulatory Commission costlessly extends the current 40-year maximum license span for a nuclear power plant to 60 years. According to the model, this would roughly double the expected present value of profits of U.S. nuclear plants and double electrical power generation over the life of the industry. With 40-year licenses, it is not optimal to repair major problems that occur after 20 years of operation. Thus, the model predicts rapid retirement of nuclear plants over the next 20 years under the current regime, whereas under 60-year licenses the number of plants would remain roughly flat for the next 20 years.

Three of the articles (Elyakime et al., Flinn, and Miller and Sieg) structurally estimate models that either explicitly or implicitly account for market interactions. Elyakime, Laffont, Loise, and Vuong extend their previous work on estimating the primitives of models of first-price sealed-bid auctions to the (common) case in which there is a second round of bargaining. Structural estimation requires that the Bayesian–Nash equilibrium strategies, partial differential equations in this case, be numerically solved. Somewhat surprisingly, they show (as in their previous work) that the underlying private valuation distributions of both buyers and sellers can be nonparametrically identified from data on bids and sellers' reservation prices. Their estimates from a parametrically specified model show that second-round bargaining provides a better fit of the data.

Flinn develops an equilibrium labor-market model in which workers are heterogeneous in their endowed productivity (and therefore "required" effort level) in a primary production sector. Effort in that sector cannot be costlessly determined, so there is an incentive for inefficient workers to shirk. As the inefficient workers are gradually discovered and dismissed, the average ability of the remaining workers rises. This dynamic selection process leads to rising wages with age in the primary sector. Flinn's results indicate that the dynamic selection process can rationalize the observed concavity of age–earnings profiles, as well as observed negative correlation between dismissals and subsequent wages.

Miller and Seig structurally estimate a model of housing consumption and male leisure based on first-order conditions using panel data from the United States and Germany. Estimating the primitives of the model enables them not only to discern whether preference structures differ significantly between the two populations (they do not) but also to determine whether international markets are fully integrated (they are not, but purchasing power parity seems to hold with regard to wages). Their ability to draw such inferences relies heavily on a complete markets characterization of the equilibrium of the economies, and their article discusses in depth the methodological issues involved in testing this characterization.

Lillard and Weiss estimate the exact reduced-form decision rule for a model of the saving behavior of retired households. They are able to obtain an exact solution because of the assumption that utility is quadratic in consumption and bequests. This allows for a consumption smoothing motive for saving but rules out a precautionary motive. Their results indicate that the bequest and consumption smoothing motives alone, combined with uncertainty about medical costs and life span, are sufficient to rationalize the observed saving of elderly married couples. Their results also indicate that an unexpected increase in Social Security benefits would lead to a proportionately large increase in saving by elderly couples because of the strength of the bequest motive.

Van den Berg and Gorter demonstrate how information on subjectively reported commuting-time-dependent reservation wages can be used to estimate the utility trade-off between wages and commuting distance in a general nonstationary job-search model. Their framework allows the recovery of the structural utility trade-off without having to estimate the full set of structural parameters. Their model also incorporates both geographic and employment mobility.

Behrman, Foster, and Rosenzweig estimate approximations to the optimal savings decision rules of farmers in rural Pakistan. The key innovation is that the seasonal timing of savings decisions is considered. In their model, farmers make savings decisions both in the planting and harvest periods of the agricultural season. They show that use of the entire season as the decision period (as in previous research) can lead to severe bias in estimating the saving–income relationship because it fails to account for the fact that much of harvest-period income is unknown at the time of the planting-period saving decision. Although harvest income shocks have a positive effect on net harvest period savings, failure to account for the seasonal timing of income completely masks this effect (i.e., total season income is unrelated to harvest-period saving).

3. AREAS FOR FUTURE RESEARCH

There are several important areas for future research related to the theme of this volume. We conclude by discussing some of these areas.

In many economic models, agents are assumed to solve (or behave as if they solve) complex optimization problems. Full-solution structural estimation requires that the econometrician must solve the mathematical representation of the same problem at each trial parameter vector during the estimation process. This can be extremely computationally burdensome. An important and active area of research is the development of fast yet accurate approximate solution methods for dynamic optimization problems (see Bellman, Kolaba, and Kotkin 1963; Geweke, Slonim, and Zarkin 1992; Keane and Wolpin 1994; Rust 1995a,b).

Another important area for research is the development of methods of structural estimation that relax assumptions about the distributions of unobservables (e.g., the shocks to preferences and technology). One approach is the semiparametric estimation of error distributions, and another approach is to use very flexible parametric distributions (e.g., scaled mixtures of normals) that can provide arbitrarily good approximations to general distributions. Applications of such approaches are still in their infancy.

The development of methods for Bayesian inference for the primitives of dynamic economic models is also of interest. This is a very difficult area because the mapping from the data to the posterior distribution of the primitive parameters is often highly nonlinear as a result of the typically highly nonlinear nature of the decision rules. In the infinite-horizon job-search case considered by Lancaster, the mapping is particularly tractable, but more general cases need to be examined. Another investigation in this area is that of Geweke and Keane (in press).

A key problem confronting the structural approach is the area of formal model selection. Classical approaches to specification testing examine the null hypothesis that a particular model is the true data-generating process. But, given the complexity of human behavior, any economic model must be a simplification of that behavior. The hope is that, although a particular model is literally false, it is nevertheless useful (say, for understanding behavior, forecasting behavior under different regimes, etc.). At the moment, formal statistical methods to choose among a set of structurally estimated economic models, all of which are assumed a priori to be false, have not been examined.

In our view, the examination of out-of-sample forecasting performance is the best available means of choosing among structurally estimated models. By out-of-sample forecast performance, we do not mean simply performance in hold-out or split-sample-type fit tests. Rather, we refer primarily to forecasting behavior in regimes not observed in the data used for estimation. At present, there are few examples in the literature in which the ability of structurally estimated models to forecast behavior under a new regime has been examined (some examples are Lumsdaine, Stock, and Wise 1992; Rust 1995a,b; Erdem and Keane 1996; Keane and Moffitt 1996). But this is largely attributable to the fact that the advent of structural estimation is rather recent, combined with the fact that in many areas a true regime change is rare.

A final important area for future research is to compare the ability of structural versus nonstructural approaches to forecast the effects of regime shifts. As we noted previously, such comparisons can only be made in instances in which the regime shift involves a change along a dimension that varies in the data so that nonstructural methods are able to generate forecasts. Nonstructural methods that could be used in such instances are (a) reduced-form representations of decision rules, (b) approximate decision rules, (c) various atheoretical approaches (e.g., vector autoregressions), or (d) the cataloging of various estimates of "treatment effects" from social and "natural" experiments followed by extrapolation to other regimes (perhaps using meta-analytic or response surface methodologies).

The desire to better forecast behavior after regime changes was a key impetus to structural estimation. Both Marschak (1952) and Lucas (1976) noted that parameters of reduced-form and statistical models, because they are functions of the primitive parameters characterizing preferences and constraints, would in general shift in a regime change that altered constraints. Thus, structurally estimated economic models—whose parameters are primitives that either remain invariant or change in a known way with a regime change—would be needed to forecast behavior in a new regime. But, given that structurally estimated economic models are not literally true, the Lucas critique provides no a priori reason to believe that structural approaches will outperform nonstructural methods for predicting effects of regime shifts. This is an empirical question on which further work is needed.

In conclusion, we believe that use of economic models to guide empirical work provides the only means to develop cumulative generalizable knowledge about economic behavior. As the articles in this volume illustrate, empirical work that is guided by theory has already led to considerable progress in applied microeconomics. The challenges confronting future work in this area are many, but we also expect that the long-term payoffs from confronting those challenges will be substantial.

[Received September 1996.]

REFERENCES

Angrist, J. (1995), "Introduction to the *JBES* Symposium on Program and Policy Evaluation," *Journal of Business & Economic Statistics*, 13, 133–136.

Bellman, R., Kalaba, R., and Kotkin, B. (1963), "Polynomial Approximation—A New Computational Technique in Dynamic Programming: Allocation Processes," *Mathematics of Computation*, 1, 155–161.

Erdem, T., and Keane, M. P. (1996), "Decision Making Under Uncertainty: Capturing Dynamic Brand Choice Processes in Turbulent Consumer Goods Markets," *Marketing Science*, 15, 1–20.

Geweke, J. F., and Keane, M. P. (in press), "Bayesian Inference for Dynamic Discrete Choice Models Without the Need for Dynamic Programming," in *Simulation Based Inference and Econometrics: Methods and Applications*, eds. R. Mariano, T. Schuerman, and M. Weeks, Cambridge, U.K.: Cambridge University Press.

Geweke, J. F., Slonim, R., and Zarkin, G. (1992), "Econometric Solution Methods for Dynamic Discrete Choice Problems," mimeo, University of Minnesota, Department of Economics.

Keane, M. P., and Moffitt, R. A. (1996), "A Structural Model of Multiple Welfare Program Participation and Labor Supply," working paper, Federal Reserve Bank of Minneapolis.

Keane, M. P., and Wolpin, K. I. (1994), "Solution and Estimation of Discrete Choice Dynamic Programming Models by Simulation and Interpolation: Monte Carlo Evidence," *Review of Economics and Statistics*, 76, 648–672.

Lucas, R. (1976), "Econometric Policy Evaluation: A Critique," *Carnegie-Rochester Conference Series on Public Policy*, 1, 19–46.

Lumsdaine, R., Stock, J., and Wise, D. (1992), "Three Models of Retirement: Computational Complexity vs. Predictive Validity," in *Topics in the Economics of Aging*, ed. D. Wise, Chicago: University of Chicago Press, pp. 21–57.

Marschak, J. (1952), "Economic Measurements for Policy and Prediction," in *Studies in Econometric Method*, eds. W. C. Hood and T. J. Koopmans, New York: Wiley, pp. 1–26.

Rust, J. (1995a), "Using Randomization to Break the Curse of Dimensionality," mimeo, University of Wisconsin, Dept. of Economics.

——— (1995b), "Structural Estimation of Markov Decision Processes," in *Handbook of Econometrics* (Vol. 4), eds. R. Engle and D. McFadden, Amsterdam: North-Holland, pp. 3081–3143.

Wolpin, K. I. (1996), "Public-Policy Uses of Discrete Choice Dynamic Programming Models," *American Economic Review*, 86, 427–432.