

THE STATA JOURNAL

Editor

H. Joseph Newton
Department of Statistics
Texas A & M University
College Station, Texas 77843
979-845-3142; FAX 979-845-3144
jnewton@stata-journal.com

Executive Editor

Nicholas J. Cox
Department of Geography
University of Durham
South Road
Durham City DH1 3LE UK
n.j.cox@stata-journal.com

Associate Editors

Christopher Baum
Boston College

Rino Bellocco
Karolinska Institutet

David Clayton
Cambridge Inst. for Medical Research

Mario A. Cleves
Univ. of Arkansas for Medical Sciences

Charles Franklin
University of Wisconsin, Madison

Joanne M. Garrett
University of North Carolina

Allan Gregory
Queen's University

James Hardin
University of South Carolina

Stephen Jenkins
University of Essex

Jens Lauritsen
Odense University Hospital

Stanley Lemeshow
Ohio State University

J. Scott Long
Indiana University

Thomas Lumley
University of Washington, Seattle

Roger Newson
King's College, London

Marcello Pagano
Harvard School of Public Health

Sophia Rabe-Hesketh
University of California, Berkeley

J. Patrick Royston
MRC Clinical Trials Unit, London

Philip Ryan
University of Adelaide

Mark E. Schaffer
Heriot-Watt University, Edinburgh

Jeroen Weesie
Utrecht University

Jeffrey Wooldridge
Michigan State University

Stata Press Production Manager

Lisa Gilmore

Copyright Statement: The Stata Journal and the contents of the supporting files (programs, datasets, and help files) are copyright © by StataCorp LP. The contents of the supporting files (programs, datasets, and help files) may be copied or reproduced by any means whatsoever, in whole or in part, as long as any copy or reproduction includes attribution to both (1) the author and (2) the Stata Journal.

The articles appearing in the Stata Journal may be copied or reproduced as printed copies, in whole or in part, as long as any copy or reproduction includes attribution to both (1) the author and (2) the Stata Journal.

Written permission must be obtained from StataCorp if you wish to make electronic copies of the insertions. This precludes placing electronic copies of the Stata Journal, in whole or in part, on publicly accessible web sites, file servers, or other locations where the copy may be accessed by anyone other than the subscriber.

Users of any of the software, ideas, data, or other materials published in the Stata Journal or the supporting files understand that such use is made without warranty of any kind, by either the Stata Journal, the author, or StataCorp. In particular, there is no warranty of fitness of purpose or merchantability, nor for special, incidental, or consequential damages such as loss of profits. The purpose of the Stata Journal is to promote free communication among Stata users.

The *Stata Journal* (ISSN 1536-867X) is a publication of Stata Press, and Stata is a registered trademark of StataCorp LP.

The Stata Journal publishes reviewed papers together with shorter notes or comments, regular columns, book reviews, and other material of interest to Stata users. Examples of the types of papers include 1) expository papers that link the use of Stata commands or programs to associated principles, such as those that will serve as tutorials for users first encountering a new field of statistics or a major new technique; 2) papers that go “beyond the Stata manual” in explaining key features or uses of Stata that are of interest to intermediate or advanced users of Stata; 3) papers that discuss new commands or Stata programs of interest either to a wide spectrum of users (e.g., in data management or graphics) or to some large segment of Stata users (e.g., in survey statistics, survival analysis, panel analysis, or limited dependent variable modeling); 4) papers analyzing the statistical properties of new or existing estimators and tests in Stata; 5) papers that could be of interest or usefulness to researchers, especially in fields that are of practical importance but are not often included in texts or other journals, such as the use of Stata in managing datasets, especially large datasets, with advice from hard-won experience; and 6) papers of interest to those teaching, including Stata with topics such as extended examples of techniques and interpretation of results, simulations of statistical concepts, and overviews of subject areas.

For more information on the Stata Journal, including information for authors, see the web page

<http://www.stata-journal.com>

Subscriptions are available from StataCorp, 4905 Lakeway Drive, College Station, Texas 77845, telephone 979-696-4600 or 800-STATA-PC, fax 979-696-4601, or online at

<http://www.stata.com/bookstore/sj.html>

Subscription rates:

Subscriptions mailed to US and Canadian addresses:

3-year subscription (includes printed and electronic copy)	\$153
2-year subscription (includes printed and electronic copy)	\$110
1-year subscription (includes printed and electronic copy)	\$ 59
1-year student subscription (includes printed and electronic copy)	\$ 35

Subscriptions mailed to other countries:

3-year subscription (includes printed and electronic copy)	\$225
2-year subscription (includes printed and electronic copy)	\$158
1-year subscription (includes printed and electronic copy)	\$ 83
1-year student subscription (includes printed and electronic copy)	\$ 59
3-year subscription (electronic only)	\$153

Back issues of the Stata Journal may be ordered online at

<http://www.stata.com/bookstore/sj.html>

The Stata Journal is published quarterly by the Stata Press, College Station, Texas, USA.

Address changes should be sent to the Stata Journal, StataCorp, 4905 Lakeway Drive, College Station TX 77845, USA, or email sj@stata.com.

Production function estimation in Stata using inputs to control for unobservables

Amil Petrin
University of Chicago
National Bureau of Economic Research

Brian P. Poi
StataCorp

James Levinsohn
University of Michigan
National Bureau of Economic Research

Abstract. A key issue in the estimation of production functions is the correlation between unobservable productivity shocks and input levels. Profit-maximizing firms respond to positive productivity shocks by expanding output, which requires additional inputs. Negative shocks lead firms to pare back output, decreasing their input usage. Olley and Pakes (1996) develop an estimator that uses investment as a proxy for these unobservable shocks. More recently, Levinsohn and Petrin (2003a) introduce an estimator that uses intermediate inputs as proxies, arguing that intermediates may respond more smoothly to productivity shocks. This paper reviews Levinsohn and Petrin’s approach and introduces a Stata command that implements it.

Keywords: st0060, levpct, production functions, productivity, endogeneity, GMM estimator

1 Introduction

A key issue in the estimation of production functions is the correlation between unobservable productivity shocks and input levels.¹ Profit-maximizing firms respond to positive productivity shocks by expanding output, which requires additional inputs. Negative shocks lead firms to pare back output, decreasing their input usage. When true, ordinary least squares (OLS) estimates of production functions are biased and, by implication, lead to biased estimates of productivity, often the relevant quantity for the estimation question.

Olley and Pakes (1996) (OP) develop an estimator that uses investment as a proxy for these unobservable shocks. More recently, Levinsohn and Petrin (2003a) (LP) point to the evidence from firm-level datasets that suggest investment is very lumpy (that is, that there are substantial adjustment costs). If this is true, the investment proxy may not smoothly respond to the productivity shock, violating the consistency condition.

¹For an overview of the history of this discussion, see Griliches and Mairesse (1998).

LP show the conditions under which intermediate inputs can also solve this simultaneity problem. Remarkably, in most applications, these inputs are not used beyond subtracting them from the gross-output number to get value added, so the approach comes at no additional cost in data or computation. LP discuss the theoretical benefits of extending the proxy choice set in this direction and provide substantial empirical evidence that these benefits are important.

One benefit is strictly data driven. It turns out that the investment proxy is only valid for plants reporting nonzero investment. (This is due to an invertibility condition described below.) Pronounced adjustment costs, which do not invalidate the use of investment as a proxy, are the likely reason that over one half of the Chilean sample used later reports zero investment. This kind of severe truncation is not unique to the Chilean data. Much of the plant-level research being conducted today is on data from countries such as India, Turkey, Columbia, Mexico, and Indonesia, and in these datasets—as is likely with others—the “zero investment” problem looms large.

Using intermediate input proxies instead of investment avoids truncating all the zero investment firms. In the data above (at least), firms almost always report positive use of intermediate inputs like electricity or materials.

To the extent that adjustment costs are an important issue, intermediate inputs may confer another benefit. If it is less costly to adjust the intermediate input, it may respond more fully to the entire productivity term than investment. For example, if adjustment costs lead to kink points in the investment demand function, plants may not respond fully to productivity shocks, and some correlation between the regressors and the error term can remain.

Another nice feature of the intermediate input is that it provides a simple link between the estimation strategy and the economic theory, primarily because intermediate inputs are not typically state variables. Levinsohn and Petrin (2003a) develop this link, showing the (mild) conditions that must hold if intermediate inputs are to be a valid proxy for the productivity shock. They suggest three specification tests for evaluating any proxy’s performance. In addition, they derive the expected directions of bias on the OLS estimates relative to LP’s intermediate input approach when simultaneity exists. Finally, LP show for the four largest Chilean manufacturing industries that significant differences between OLS and Levinsohn–Petrin exist that are exactly consistent with simultaneity.²

This paper reviews Levinsohn and Petrin’s approach and provides a Stata command that implements it.

2 Productivity estimation

In this section, we give an overview with an emphasis on the mechanics of the estimator. A more detailed exposition can be found in Levinsohn and Petrin (2003a).

²This finding has been reported for a number of other manufacturing surveys, including surveys from Columbia and India.

For the purposes of this note, the production technology is assumed to be Cobb–Douglas

$$y_t = \beta_0 + \beta_l l_t + \beta_k k_t + \beta_m m_t + \omega_t + \eta_t \quad (1)$$

where y_t is the logarithm of the firm’s output, most often measured as gross revenue or value added; l_t and m_t are the logarithm of the freely variable inputs labor and the intermediate input; and k_t is the logarithm of the state variable capital.³ For ease of exposition, we include only two freely variable inputs, though the command accompanying this article allows for an arbitrary number of them.

The error has two components: the transmitted productivity component given as ω_t and η_t , an error term that is uncorrelated with input choices. The key difference between ω_t and η_t is that the former is a state variable and, hence, impacts the firm’s decision rules. It is not observed by the econometrician, and it can impact the choices of inputs, leading to the well-known simultaneity problem in production function estimation. Estimators ignoring this correlation between inputs and this unobservable factor (like OLS) will yield inconsistent results.

Demand for the intermediate input m_t is assumed to depend on the firm’s state variables k_t and ω_t :

$$m_t = m_t(k_t, \omega_t)$$

Making mild assumptions about the firm’s production technology, Levinsohn and Petrin (2003a, appendix A) show that the demand function is monotonically increasing in ω_t . This allows inversion of the intermediate demand function, so ω_t can be written as a function of k_t and m_t :

$$\omega_t = \omega_t(k_t, m_t)$$

The unobservable productivity term is now expressed solely as a function of two observed inputs.

A final identification restriction follows Olley and Pakes (1996). LP assume that productivity is governed by a first-order Markov process

$$\omega_t = E[\omega_t | \omega_{t-1}] + \xi_t$$

where ξ_t is an innovation to productivity that is uncorrelated with k_t , but not necessarily with l_t ; this is part of the source of the simultaneity problem.

First, we discuss estimation when the dependent variable is value added. Then, we turn to using output (or gross revenue) as the dependent variable.

2.1 Estimation in the value-added case

Letting v_t represent value added—gross-output net of intermediate inputs—we can write the production function as

³The approach extends immediately to other forms of the production technology. For example, with an appropriate definition of variables, trans-log (and higher order) production functions can be estimated with this Stata routine.

$$\begin{aligned} v_t &= \beta_0 + \beta_l l_t + \beta_k k_t + \omega_t + \eta_t \\ &= \beta_l l_t + \phi_t(k_t, m_t) + \eta_t \end{aligned}$$

where

$$\phi_t(k_t, m_t) = \beta_0 + \beta_k k_t + \omega_t(k_t, m_t)$$

Substituting a third-order polynomial approximation in k_t and m_t in place of $\phi_t(k_t, m_t)$, makes it possible to consistently estimate parameters of the value-added equation using OLS as

$$v_t = \delta_0 + \beta_l l_t + \sum_{i=0}^3 \sum_{j=0}^{3-i} \delta_{ij} k_t^i m_t^j + \eta_t$$

where β_0 is not separately identified from the intercept of $\phi_t(k_t, m_t)$.⁴ This completes the first stage of the estimation routine from Levinsohn and Petrin (2003a), from which an estimate of β_l and an estimate of ϕ_t (up to the intercept) are available.

The second stage of the routine identifies the coefficient β_k . It begins by computing the estimated value for ϕ_t using

$$\begin{aligned} \widehat{\phi}_t &= \widehat{v}_t - \widehat{\beta}_l l_t \\ &= \widehat{\delta}_0 + \sum_{i=0}^3 \sum_{j=0}^{3-i} \widehat{\delta}_{ij} k_t^i m_t^j - \widehat{\beta}_l l_t \end{aligned}$$

For any candidate value β_k^* , we can compute (up to a scalar constant) a prediction for ω_t for all periods t using

$$\widehat{\omega}_t = \widehat{\phi}_t - \beta_k^* k_t$$

Using these values, a consistent (nonparametric) approximation to $E[\omega_t | \omega_{t-1}]$ is given by the predicted values from the regression

$$\widehat{\omega}_t = \gamma_0 + \gamma_1 \omega_{t-1} + \gamma_2 \omega_{t-1}^2 + \gamma_3 \omega_{t-1}^3 + \epsilon_t$$

which LP call $E[\widehat{\omega}_t | \widehat{\omega}_{t-1}]$.

Given $\widehat{\beta}_l$, β_k^* , and $E[\widehat{\omega}_t | \widehat{\omega}_{t-1}]$, LP write the sample residual of the production function as

$$\widehat{\eta}_t + \xi_t = v_t - \widehat{\beta}_l l_t - \beta_k^* k_t - E[\widehat{\omega}_t | \widehat{\omega}_{t-1}]$$

Our estimate $\widehat{\beta}_k$ of β_k is defined as the solution to

$$\min_{\beta_k^*} \sum_t (v_t - \widehat{\beta}_l l_t - \beta_k^* k_t - E[\widehat{\omega}_t | \widehat{\omega}_{t-1}])^2$$

The Stata command accompanying this note uses a golden section search algorithm to minimize that function. A bootstrap approach (discussed shortly) is used to construct standard errors for $\widehat{\beta}_l$ and $\widehat{\beta}_k$. We now turn to point estimation when the dependent variable is output (or gross revenue).

⁴Another restriction is necessary to separately identify β_0 from the intercept of $\phi_t(k_t, m_t)$.

2.2 Estimation in the gross revenue case

Letting y_t denote revenue, the production function is given as

$$\begin{aligned} y_t &= \beta_0 + \beta_l l_t + \beta_k k_t + \beta_m m_t + \omega_t + \eta_t \\ &= \beta_l l_t + \phi_t(k_t, m_t) + \eta_t \end{aligned}$$

where now

$$\phi_t(k_t, m_t) = \beta_0 + \beta_k k_t + \beta_m m_t + \omega_t(k_t, m_t)$$

Estimation of $\widehat{\beta}_l$ proceeds exactly as before, using OLS with a third-order polynomial approximation in k_t and m_t in place of $\phi_t(k_t, m_t)$.

The first part of the second stage is also similar to the value-added case. For any candidate values β_k^* and β_m^* (for β_k and β_m), we estimate $\widehat{\omega}_t$ using

$$\widehat{\omega}_t = \widehat{\phi}_t - \beta_k^* k_t - \beta_m^* m_t$$

Using the ω_t 's for all t , we estimate $E[\widehat{\omega}_t | \omega_{t-1}]$ as before. Then, the residual for (β_k^*, β_m^*) is computed as

$$\eta_t + \widehat{\xi}_t = y_t - \widehat{\beta}_l l_t - \beta_k^* k_t - \beta_m^* m_t - E[\widehat{\omega}_t | \omega_{t-1}]$$

This residual must interact with at least two instruments to identify both β_k and β_m . Similar to the value-added case, if period t 's capital stock is determined by the previous period's investment decisions, it does not respond to shocks to this period's productivity innovation term ξ_t , providing the moment condition

$$E[\eta_t + \xi_t | k_t] = 0$$

which is implicitly imposed in the objective function from (1). An additional moment condition is needed to identify β_m separately from β_k . LP use the fact that the previous period's level of material usage m_t is uncorrelated with this period's error, giving us the moment condition

$$E[\eta_t + \xi_t | m_{t-1}] = 0$$

Thus, with $\mathbf{Z}_t \equiv (k_t, m_{t-1})$, one candidate estimator solves

$$\min_{(\beta_k^*, \beta_m^*)} \sum_h \left\{ \sum_t (\eta_t + \widehat{\xi}_t) Z_{ht} \right\}^2$$

with h indexing the elements of \mathbf{Z}_t .

Additional overidentification conditions are given by

$$E[\eta_t + \xi_t | l_{t-1}] = 0, \quad E[\eta_t + \xi_t | m_{t-2}] = 0, \quad \text{and} \quad E[\eta_t + \xi_t | k_{t-1}] = 0$$

These can be used to improve efficiency and test the specification. Here one redefines $\mathbf{Z}_t \equiv (k_t, m_{t-1}, l_{t-1}, m_{t-2}, k_{t-1})$. $\widehat{\beta}_k$ and $\widehat{\beta}_m$ are then defined as the solution to

$$\min_{(\beta_k^*, \beta_m^*)} \sum_h \left\{ \sum_t (\eta_t + \widehat{\xi}_t) Z_{ht} \right\}^2$$

Our Stata implementation provides two methods for solving the GMM minimization problem. The default behavior is to use Stata's `nl` command, which is based on Newton's method. Alternatively, a two-dimensional grid search can be requested. Candidate values for β_k^* and β_m^* from 0.01 to 0.99, in increments of 0.01, are used. Although much slower than `nl`, the grid search is handy for confirming that `nl` has found the global minimum of the objective function. Moreover, if there is insufficient variation in the capital and proxy variables, `nl` may have difficulty solving the minimization problem; in these cases, we can instead use the grid search.

2.3 Standard errors

The estimators developed above involve two main stages of estimation. In each of these stages, a number of preliminary estimators are used. The covariance matrix of the final parameters must account for the sampling variation introduced by all of the estimators used in the two stages. Although deriving an analytic covariance matrix may be feasible, this calculation is not trivial. Instead, LP substitute computational power for analytic difficulties, employing the bootstrap to estimate standard errors.

Because LP use panel data, they sample with replacement from firms, using the entire time series of observations for that firm in the bootstrapped sample when the firm's id number is randomly drawn. The variation in the point estimates across the bootstrapped samples provides an estimate for the standard errors of the original point estimates. In Stata, we accomplish the bootstrap sampling using the `cluster(varname)` option with the `bsample` command.

Bootstrapping, when overidentifying restrictions are imposed, is slightly different. The sample moments computed using the original dataset will, in general, not equal zero, even though the population moments do (by assumption).⁵ As Horowitz (2001) and others have noted, this means that, for each of the bootstrapped samples, one must "recenter" the moment conditions by subtracting the values of the sample moments calculated using the original dataset (at the minimum). Our Stata implementation accomplishes this by first performing the estimation on the original dataset. We then store the values of the sample moments using a series of global macros. Their value is then subtracted from the bootstrapped sample's moments when minimizing the objective function for that bootstrapped sample. This restores the consistency of the bootstrap approach in the construction of standard errors.

3 Stata implementation

3.1 Syntax

```
levpet depvar [if exp] [in range], free(varlist) proxy(varlist)
      capital(varname) [[valueadded | revenue] justid grid i(varname)
      t(varname) reps(#) level(#)]
```

⁵See Horowitz (2001) for an overview of the bootstrap and a discussion of the necessity of recentering.

Syntax for predict

`predict` [*type*] *newvarname* [*if exp*] [*in range*], *omega*

3.2 Options

`free`(*varlist*) specifies the freely variable inputs, excluding the one used as the proxy variable.

`proxy`(*varlist*) specifies the proxy variables, which are typically electricity, materials, or fuels. If the dependent variable is value added, one or two variables can be specified. If the dependent variable is revenue, only one can be specified.

`capital`(*varname*) specifies the capital variable.

`valueadded` indicates that the dependent variable represents value added and that the least-squares estimator be used. This is the default.

`revenue` indicates that the dependent variable represents gross revenue and that the GMM estimator will be used.

`justid` requests that the GMM estimator use only present-period capital and the first lag of the proxy variable as instruments. The default is to include lagged labor, lagged capital, and the second lag of the proxy variable as instruments as well. This option can only be used with `revenue`.

`grid` requests that the GMM estimator use a grid search to minimize the criterion function with respect to the coefficients on capital and the intermediate input. The default is to use Stata's `nl` command. This option can only be used with `revenue`.

`i`(*varname*) specifies the variable that contains the unit to which the observation belongs. You can specify the `i()` option the first time you estimate, or you can use the `iis` command to set `i()` beforehand. Note that it is not necessary to specify `i()` if the data has been previously `tsset` or if `iis` has been previously specified. In these cases, the group variable is taken from the previous setting. See [XT] `xt`.

`t`(*varname*) specifies the variable that contains the time at which the observation was made. You can specify the `t()` option the first time you estimate, or you can use the `tis` command to set `t()` beforehand. Note that it is not necessary to specify `t()` if the data has been previously `tsset` or if `tis` has been previously specified. In these cases, the time variable is taken from the previous setting. See [XT] `xt`.

`reps`(*#*) specifies the number of bootstrap replications to be performed. The default is 50.

`level`(*#*) specifies the confidence level, as a percentage, for confidence intervals. The default is `level(95)` or as set by `set level`; see [U] **23.6 Specifying the width of confidence intervals**.

Options for predict

`omega` requests the predicted levels of productivity, where for the value-added case

$$\widehat{\omega}_t = \exp(v_t - \widehat{\beta}_l l_t - \widehat{\beta}_k k_t)$$

and for the gross revenue case

$$\widehat{\omega}_t = \exp(y_t - \widehat{\beta}_l l_t - \widehat{\beta}_k k_t - \widehat{\beta}_m m_t)$$

`Predict` assumes that the production function inputs are in log levels and adjusts ω_t accordingly. If there is more than one freely variable input, these formulas are modified accordingly.

3.3 Remarks

The `levpet` command implements the Levinsohn–Petrin estimator, as discussed in the previous section. The command works with versions 7 and higher of Stata. The dialog box requires Stata 8 and can be invoked by typing

```
. db levpct
```

4 Example

Here, we illustrate the usage of `levpet` using a dataset consisting of Chilean apparel firms from 1987 through 1996. We have data on value added, as well as firms' usage levels of blue and white collar labor, electricity, and capital. We treat blue and white collar labor as freely variable inputs, and we use electricity as the proxy variable. In the examples below, the prefix `ln` on a variable name indicates the natural log.

We consider the case in which the dependent variable represents value added. In Stata, we type

```
. tsset ppn year
. levpct lnva, free(lnb lnw) proxy(lne) capital(lnk) valueadded reps(250)
```

The output looks similar to most of Stata's `xt` commands:

(Continued on next page)

At the bottom of table 1, we also report the sum of the coefficients for each estimator; constant returns to scale correspond to a sum of one. Both OLS and LP imply increasing returns for this industry, though in the case of LP, we cannot reject the null that the sum is one at the 5% significance level.

5 Conclusion

When estimating production functions, we must account for the correlation between input levels and productivity. Profit-maximizing firms respond to increases in productivity by increasing their usage of factor inputs. Methods that ignore this endogeneity, such as OLS and the fixed-effects estimator, will provide inconsistent estimates of the parameters of the production function.

Building on the work of Olley and Pakes (1996), Levinsohn and Petrin (2003a) develop an estimator that utilizes intermediate inputs to proxy for the unobservable productivity term. Most plant-level datasets include data on the usage of intermediate inputs such as energy and materials, so Levinsohn and Petrin's estimator does not suffer from the truncation bias induced by Olley and Pakes' estimator, which requires firms to have nonzero levels of investment.

In this paper, we have introduced the Stata command `levpet` to implement this estimator. We hope that its simple syntax will motivate people to consider it as a better alternative to estimators that ignore endogeneity issues.

6 Saved results

`levpet` saves in `e()`:

Scalars

<code>e(N)</code>	number of observations
<code>e(waldcrs)</code>	Wald test of constant returns

Macros

<code>e(cmd)</code>	<code>levpet</code>
<code>e(predict)</code>	program used to implement <code>predict</code>
<code>e(model)</code>	<code>value added</code> or <code>revenue</code> ; model used

Matrices

<code>e(b)</code>	coefficient vector
<code>e(V)</code>	variance-covariance matrix

Functions

<code>e(sample)</code>	marks estimation sample
------------------------	-------------------------

7 References

- Griliches, Z. and J. Mareisse. 1998. Production functions: The search for identification. In *Econometrics and Economic Theory in the Twentieth Century: The Ragnar Frisch Centennial Symposium*, 169–203. Cambridge: Cambridge University Press.
- Horowitz, J. L. 2001. The bootstrap. In *Handbook of Econometrics*, ed. J. J. Heckman and E. Leamer, vol. 5, 3159–3228. Amsterdam: North-Holland.
- Levinsohn, J. and A. Petrin. 2003a. Estimating production functions using inputs to control for unobservables. *Review of Economic Studies* 70(2): 317–342.
- . 2003b. On the micro-foundations of productivity growth. Mimeo: University of Chicago.
- Olley, G. S. and A. Pakes. 1996. The dynamics of productivity in the telecommunications equipment industry. *Econometrica* 64: 1263–1297.

About the Authors

James Levinsohn is Professor of Economics and Public Policy at the University of Michigan and a member of the National Bureau of Economic Research.

Amil Petrin is Associate Professor of Economics and Statistics at the University of Chicago Graduate School of Business and a member of the National Bureau of Economic Research.

Brian Poi is an economist and Assistant Director of Technical Services at StataCorp.