

American Economic Association

Understanding Productivity: Lessons from Longitudinal Microdata

Author(s): Eric J. Bartelsman and Mark Doms

Source: *Journal of Economic Literature*, Vol. 38, No. 3 (Sep., 2000), pp. 569-594

Published by: American Economic Association

Stable URL: <http://www.jstor.org/stable/2565420>

Accessed: 17/03/2009 13:27

Your use of the JSTOR archive indicates your acceptance of JSTOR's Terms and Conditions of Use, available at <http://www.jstor.org/page/info/about/policies/terms.jsp>. JSTOR's Terms and Conditions of Use provides, in part, that unless you have obtained prior permission, you may not download an entire issue of a journal or multiple copies of articles, and you may use content in the JSTOR archive only for your personal, non-commercial use.

Please contact the publisher regarding any further use of this work. Publisher contact information may be obtained at <http://www.jstor.org/action/showPublisher?publisherCode=aea>.

Each copy of any part of a JSTOR transmission must contain the same copyright notice that appears on the screen or printed page of such transmission.

JSTOR is a not-for-profit organization founded in 1995 to build trusted digital archives for scholarship. We work with the scholarly community to preserve their work and the materials they rely upon, and to build a common research platform that promotes the discovery and use of these resources. For more information about JSTOR, please contact support@jstor.org.



American Economic Association is collaborating with JSTOR to digitize, preserve and extend access to *Journal of Economic Literature*.

<http://www.jstor.org>

Understanding Productivity: Lessons from Longitudinal Microdata

ERIC J. BARTELSMAN *and* MARK DOMS¹

1. Introduction

RECENT YEARS HAVE seen a dramatic increase in studies on productivity that use longitudinal micro-level data sets (LMDs) which follow large numbers of establishments or firms over time. The popularity of this emerging research can be ascribed, in part, to increased availability of micro-level data, to the development of a rich theoretical microeconomic foundation, and to the displeasure with the concept of the aggregate production function. But the most important impetus has been the host of interesting questions that can be addressed effectively only with microdata. Although many productivity studies have used microdata, this review focuses on productivity studies that use LMDs, in particular the Longitudinal Research Database (LRD), which is a large panel data set of U.S. manufacturing plants developed by the U.S. Bu-

reau of the Census.² Because of the recent volume of research using the LRD and the new directions into which the research is blossoming, the time has come to step back and examine what is now known and to contemplate directions for future research.

This review complements two articles that recently appeared in this journal. First, the article by Richard Caves (1998) describes how research using LMDs to document firm and establishment demographics provides new evidence pertaining to theories of industrial organization. Next, James Tybout (2000) discusses what has been gleaned from similar microdata sets in developing countries. The literature he reviews focuses on how conditions in developing countries, such as protected markets and imperfect competition, affect productivity dispersion and productivity growth.

The literature discussed in this review is concerned with issues directly related to productivity in industrialized countries. The research papers that use LMDs can be roughly divided into two groups: those that document and describe productivity, and those that examine the

¹ Bartelsman: Free University, Amsterdam. Doms: Federal Reserve Board of Governors. The authors would like to thank the editors and referees of the *JEL* for their editorial guidance and support. Further thanks go to Martin Baily, Bert Balk, Frank den Butter, Carol Corrado, John Haltiwanger, Robert McGuckin, Joyce Zickler, and numerous others. Special thanks to the staff at the Center for Economic Studies at the US Bureau of the Census for providing a stimulating research environment, and to Ellen Dykes at the Federal Reserve Board for careful reading and editing. The views presented are solely those of the authors and not necessarily those of the Federal Reserve Board or its staff.

² Many other countries have developed data sets like the LRD. However, productivity research using the LRD is more mature. Results from other data sets will be discussed as needed.

factors behind productivity growth. The first group of papers documents the cross-sectional distributions of productivity and the evolution of industry productivity growth. These papers present a myriad of stylized facts on the dispersion of productivity across firms and establishments, the uniformity of changes in productivity, the persistence of productivity differentials, the consequences of entry and exit, and the importance of changes in resource reallocation across firms to aggregate productivity growth.

The second group of papers covered in this review attacks one of the most fundamental questions in productivity analysis: What are the factors underlying productivity growth? Some of the factors that have recently been examined include managerial ability, technology, human capital, and regulation. Although these factors are all thought to be important, not much is known about their relative importance or about the way they interact. Part of this literature has documented the correlation between productivity and variables thought to influence it. The more advanced literature in this field goes a step further and addresses the difficult question of causality. For instance, differences in technology use among firms are seen to be correlated with productivity differences. This finding leads to the more interesting question of where the technology comes from and why the use of technology varies greatly across firms.

One of the reasons that the use of LMDs for productivity research is becoming so popular is that they fill a void between two of the main subfields within the productivity literature. As a gross generalization, one of the subfields is research that focuses on growth accounting and the estimation of factor demands using aggregate and sectoral data. Some of this work was reviewed

by M. Ishaq Nadiri (1970) in this journal. A second subfield is research that examines the factors underlying changes in productivity at the firm level, namely evolutionary models of productivity growth, reviewed by Richard Nelson (1981) in this journal. Until recently, the links between these two groups have been almost nonexistent, a situation that is surprising as well as disappointing because both have the objective to gain better insight into the sources of productivity growth (or the changes in the pace of productivity growth). The evolutionary literature recognizes the large amount of heterogeneity across firms regarding their productivity and seeks to explore the factors behind this heterogeneity within the framework of firm behavior. Given the great heterogeneity that exists across firms within industries, the evolutionary literature eliminates the centerpiece of the growth-accounting literature, the representative firm. However, the results from the evolutionary literature are difficult to generalize because they are often based on anecdotes, single industries, or unrepresentative samples. Therefore, results from the evolutionary side can only be suggestive of trends in aggregate productivity growth.

A simple story line combining the complementary approaches is emerging from the reviewed literature. It provides an understanding of how changes in policy or environmental factors affect aggregate productivity, even though it is not able to pinpoint the truly exogenous causes of productivity. Figure 1 provides a simple schematic of the story.

First, innovative activity generates technology needed for production. A process of diffusion and adoption of technology determines productivity at each firm or establishment, generating a cross-sectional distribution of productivity.




Figure 1.

Next, interactions between producers in the market determine market shares. Aggregate productivity can be computed as the share-weighted average of individual productivity. This schematic aids in interpreting the stylized facts on productivity dispersion and the decomposition of aggregate productivity reviewed in Section 3.

The next part of the story line is the impact that changes in certain factors have on the processes in the schematic. For example, R&D subsidies may increase innovative activity, but they may also alter the strategic interactions between firms that determine market shares. Trade liberalization may speed up diffusion of technology and increase volatility of market shares. Innovations in financial markets that allowed for management buyouts may alter the incentives of firms to adopt certain technology. Changes in antitrust enforcement or deregulation of markets change interactions in the market, reduce entry and exit barriers, and alter incentives to invest in new production capacity. In Section 4, examples of research on the impact of these factors on productivity will be reviewed.

What have studies using data sets like the LRD taught us about productivity? First, the amount of productivity dispersion is extremely large—some firms are substantially more productive than others. Second, highly productive firms today are more than likely to be highly productive firms tomorrow, although there is a fair amount of change in the productivity distribution. Third, a large

portion of aggregate productivity growth is attributable to resource reallocation. The manufacturing sector is characterized by large shifts in employment and output across establishments every year—the aggregate data belie the tremendous amount of turmoil underneath. This turmoil is a major force contributing to productivity growth, resurrecting the Schumpeterian idea of creative-destruction. Fourth, quantifying the importance of various factors behind productivity growth, such as changes in the regulatory environment or changes in technology, is a difficult task and has been only partially successful. Nonetheless, some useful lessons have been learned. In terms of the regulatory environment, any regulations that inhibit resource reallocation can have detrimental effects on productivity growth. Regarding the effect of technology on productivity, it is now known that documenting the correlation between a factor of production, such as computers, and productivity is not enough to understand causal mechanisms. Use of computers also is related to other variables correlated with productivity, such as human capital and managerial ability.

Although this review is about what has been learned about productivity growth using microdata, it is worth mentioning some of the other fields of economics that have benefited from data sets like the LRD. In macroeconomics, the study of productivity moved to the forefront with the advent of real business cycle models in which macroeconomic fluctuations are assumed to be driven by

technology shocks as measured by cyclical movements in total factor productivity (TFP). Much of what has been learned about productivity from the analysis of microdata is currently being used to improve the understanding of the origins and propagation of business cycle fluctuations and sheds new light over the debate on procyclical productivity residuals.

In labor economics, productivity has come to the forefront in discussions about skill biases of technical change, wage dispersion, and cyclical fluctuations in employment gross flows. For example, recent work has used microdata to identify more accurately the relationships between workers' skills and technology usage, testing a necessary condition for the skill-biased technical change hypothesis.

In international trade, the effects of increased trade and competition on productivity have been widely speculated. Using LMDs that contain information on exports, several authors have addressed the question of whether or not exporting leads to increased productivity because of exposure to foreign markets. Other authors are examining the question of whether domestic producers increase their productivity when faced with foreign competition. The role of trade in the success and failure of firms in developing countries is reviewed by Tybout (2000).

The field of industrial organization has contributed much to, and learned from, research using LMDs, as described by Caves (1998). Models of industry evolution are driven to a great extent by differences in productivity, and explanations for dispersion of productivity have been forthcoming from advances in IO theory. Early theoretical work on micro-level heterogeneity played a role in stimulating the empirical analysis of microdata. Robert Lucas

(1978) and Boyan Jovanovic (1982) developed frameworks in which industries are not composed of representative firms or plants but of firms that have different efficiency endowments. Other contributions to the modeling of industry dynamics include Val Lambson (1991), and Richard Ericson and Ariel Pakes (1995). These models have been used as theoretical underpinnings for analyzing productivity growth and within-industry dynamics. For example, the approach used by Steven Olley and Pakes (1996) deals explicitly with endogenous exit behavior, as well as with simultaneity of productivity and factor demand. LMDs have been used to evaluate these models and to provide a framework in which to think about heterogeneity. Another area of industrial organization that has greatly benefited from LMDs is the literature on mergers and acquisitions. The competing theories of merger behavior can be tested empirically using LMDs to examine the productivity of establishments before and after changes in ownership in order to see to what extent productivity improves.

The outline of this review is as follows. The next section describes the data sets that have been used in the work under review, particularly the LRD. Attention is also given to the measurement of labor and total factor productivity. The third section reviews the papers that document the components in the framework mentioned above: the distribution of productivity, the changes in productivity at the micro level, and the way these changes are related to the level of and changes in market shares. The fourth section reviews work that has examined the factors, such as technology, human capital, and international exposure to competition, underlying the patterns. Although the contributions of the reviewed work are significant, a

host of issues in productivity research—problems that are especially relevant to microdata—remain unresolved. These issues include the way to measure productive inputs and outputs, the appropriate frequency of the data, and the applicability or relevance of specific microresults to understanding aggregate phenomena. The fifth section provides suggestions for the direction of future research.

2. Data and Methods

This survey reviews what has recently been learned about productivity growth from the use of large panel data sets, where the unit of analysis is the establishment or the firm. Some of these data sets are maintained by government agencies, and researchers' access to them has greatly increased since the 1980s.³ The emergence of LMD research has been driven by greater access to government databases and the increase in computer power for processing the data. Zvi Griliches and Vidar Ringstad (1971) pioneered the use of LMDs in their study of scale elasticities, which used the Norwegian Census of Manufactures. Since that time, many other countries have granted access to their data.

In the United States, the LRD has been used extensively in productivity research over the past decade. It is a collection of two surveys at the manufacturing establishment level: the Census of Manufacturers (a manufacturing universe survey with more than 300,000 records, which is conducted every five years) and the Annual Survey of Manufacturers (annual samples of the manu-

facturing universe containing data on 55,000 to 75,000 establishments). These two surveys collect basic information on inputs and outputs and on firm ownership and location. The establishment data from both surveys have been linked, starting in 1972 and at present continuing through 1997, allowing the tracking of individual establishment performance over time (the earliest description of the LRD is given by Robert McGuckin and George Pascoe 1988). Many other countries have developed similar data sets, including Canada, France, Denmark, Norway, Sweden, the Netherlands, and Israel. Although these data sets are being used more intensively to address a wide set of questions, data access can still be problematic. For instance, researchers must use the LRD only at secured sites (the Census Bureau's headquarters in Maryland or one of five satellite sites throughout the country). Also, the fees to use the data run several thousand dollars a month.

The reasons that the LRD and other longitudinal data sets have recently been so popular in productivity analysis is that they contain large and statistically representative samples and they allow the linking of data from supplementary surveys.

Large Samples. The advantage of the LMDs having large, statistically representative samples is important given the tremendous heterogeneity of business populations, both in size and performance. Also, the size distribution of establishments is skewed, so obtaining representative samples can be as difficult as it is important.

The other obvious reason that the large samples of the LMDs are useful is that they allow one to control for a greater number of factors in obtaining more precise parameter estimates. In much of the work reviewed in this paper,

³ Many other microdata sets are used in productivity analysis, with information on firms in industries such as airlines, banking, health care, and utilities. This review focuses primarily on results from data sets that are more comprehensive in their industry coverage.

variables such as size, detailed industry, age, and region are controlled for in regressions. Unlike research that relies on aggregate data, research that uses LMDs exploits the variation within industries to identify the parameters of interest, instead of relying on cross-industry variation.

Time Series. The second advantage of LMDs is that they track the performance of individual units over time. The distribution of both productivity levels and growth rates is quite wide (see below). The time series permits the controlling of firm and plant fixed effects, which can be important in determining the relationships between a set of variables and productivity performance. The time series data can also be used to assess the importance of turnover in industries, for example, the importance of entry and exit and of resource allocation across establishments.

Ability to Link Data from Other Surveys. A third advantage of LMDs has been the ability to link to data from other surveys. The LRD contains variables of a general variety: total employment, wages and salaries, value of shipments, cost of materials, and so forth. Additionally, establishments in the LRD have been linked to other survey data that are related to productivity, such as the use of advanced manufacturing technology (Tim Dunne 1994), foreign ownership (Doms and Bradford Jensen 1998), and detailed worker characteristics (Doms, Dunne, and Kenneth Troske 1997), to name a few. By combining survey data at the establishment or firm level, one can address a richer set of issues than by examining each data set separately. Without linking capabilities, grouped data often need to be merged with the microdata, on the basis of a more aggregate characteristic, such as region or industry. Merging with aggregate data may result in downward bi-

ased standard errors of estimates owing to intragroup error-term correlations (see Brent Moulton 1990).

Measuring Productivity with LMDs. The availability of LMDs has not substantially altered the existing methods used for measuring productivity, although the availability of detailed microdata has stimulated development of some innovative solutions to old empirical problems. An overview of methodologies used in the research reviewed in this paper will aid in interpreting results. Readers who are familiar with total factor productivity (TFP) measurement issues and are interested mainly in the story line may safely skip ahead.

The main choice that researchers make is whether to analyze labor productivity or TFP. Next, they must choose whether the output concept is in physical terms, in deflated gross production, or in (double) deflated value added. Among methods of calculating TFP, choices can be made among index number approaches, econometric estimation of cost or production functions, or nonparametric methods, such as data envelopment analysis. If one is interested only in productivity measures *per se*, it is best not to take a dogmatic stance on methodology but rather to explore the sensitivity of productivity measures to variations in methodology. In assessing the effect of certain factors on productivity, one must take care that the chosen empirical framework allows proper identification of the relationship.

This review defines productivity as a measure of TFP, unless explicit mention is made of the fact that labor productivity is being discussed. TFP can be thought of in a manner analogous to labor productivity, namely as a ratio of output to input. A simple measure of the level of labor productivity is given by:

$$LP_{it} = (y_{it}/l_{it})$$

where LP shows the output, y , for firm i , period t , per unit of labor input, l . Numerous complications arise concerning measurement of outputs and inputs or aggregation of inputs across types, across agents, and over time. For example, labor input may be measured in employees, in hours worked, or in quality-adjusted hours worked.

Although TFP is an analogous concept, in a general multi-output multi-input setting, the level of productivity cannot be measured in any meaningful units. Instead, productivity changes of a firm over time or productivity movements relative to other firms need to be measured by use of index numbers. The equation

$$\pi_{it}/\pi_{it-1} = (y_{it}/x_{it})/(y_{it-1}/x_{it-1})$$

shows, for firm i , productivity (π) change between period t and period $t-1$, where y is an index of output quantity and x is an index of aggregate inputs.

Likewise, the equation

$$\pi_i/\pi_j = (y_i/x_i)/(y_j/x_j)$$

shows productivity of firm i relative to that of firm j , in a given period.

The choice between TFP or labor productivity measurement is fundamental because increases in labor productivity can result from increases in the capital-labor ratio without changes in underlying technology. TFP provides more information about changes in technology than does labor productivity and is the preferred concept despite problems arising from the measurement of capital service flows. For certain limited purposes, however, labor productivity may be an appropriate concept. For welfare comparisons, value added per capita is informative. At a more disaggregated level, labor productivity allows one, for example, to compare unit labor costs of production across establishments.

When establishments are known to have the same capital, labor productivity obviously suffices. Heterogeneity in labor productivity has been found to be accompanied by similar heterogeneity in TFP in the reviewed research where both concepts are measured.

The choice of output measure is often dictated by the available data. Where possible, physical output with unchanging quality is the best measure. The LRD has physical output measures for a small subset of products produced by a subset of establishments. However, these data are not corrected for quality changes over time, nor are they necessarily comparable across establishments. In general, researchers rely on deflating nominal variables at the sectoral level (all establishments in an industry use the same deflators). Most studies using the LRD use production and materials deflators at the four-digit SIC level available from a data set maintained by Bartelsman and Wayne Gray (1996). Using deflated production to measure productivity has one drawback, which is the same whether applied at the micro level or at the sectoral or aggregate level: Any quality improvement in output that is not reflected in the deflator will result in a downward bias in productivity.

The unavailability of micro-level prices, however, introduces other problems that will be briefly mentioned here. Applying quality-adjusted industry deflators to micro-level data would be acceptable under perfect competition because the price per unit of quality-adjusted output would be the same across firms. However, the persistent dispersion of productivity or costs across firms, which is discussed later, calls into question the empirical relevance of perfect competition. Under monopolistically competitive markets with differentiated products, prices may

differ across micro units. In this case, assuming constant prices implies that an establishment with higher-than-average prices will mistakenly be assigned higher productivity. Disentangling price and quantity movements, even if micro-level market prices were available, requires knowledge of the demand side for the differentiated goods. Most of the papers mentioned in this review abstract from the above-mentioned problem. A notable exception worth following is the paper by Tor Klette and Griliches (1996) who look specifically at the issue of productivity and product demand in an environment with unobserved output prices. They find that comparing changes in a firm's share of industry revenue with its share of inputs conveys information regarding process or product innovation.

Another choice facing researchers is the use of gross production or value added as the output concept. Value added may be more useful for making welfare statements at an aggregate level but less useful for understanding sources of productivity growth. Shifts in use of intermediate inputs relative to capital and labor over time may create biases in productivity measured with value-added output. As noted in Martin Baily (1986), the more disaggregated the data, the greater the advantage of using gross production for productivity measurement. Net-output or Domar aggregation of TFP measured with gross production will avoid double-counting problems (see an example in Bartelsman and Gray 1996).

Measures of TFP can be calculated using a wide variety of methods, and the comparative advantage of each measure depends on the particular question at hand and the particular constellation of available data. A general decomposition of an index of TFP change, derived by Bert Balk (1998), allows computation

of the contributions from technological change, changes in technical and allocative efficiency, effects of nonmarginal cost pricing, and effects of non-constant returns to scale. With k inputs and m outputs, the k -vector of input prices and quantities and the m -vector of output prices and quantities at time 0 and 1 can be used to compute TFP change:

$$\frac{\pi^1}{\pi^0} = \frac{\frac{p^1 y^1}{p^0 y^0} / P(y^1, y^0, p^1, p^0)}{\frac{w^1 x^1}{w^0 x^0} / W(x^1, x^0, w^1, w^0)}.$$

This equation states that the change in TFP equals deflated change in revenues divided by deflated change in costs, where the functions P and W represent a price index for deflation.

The Solow residual is a special case of this ratio and measures technological change if certain "standard" assumptions hold. These assumptions are that factors of production are utilized in a technically efficient manner and are paid their marginal products, and that outputs are produced under constant returns to scale and priced at marginal cost. The equation

$$\hat{\pi}_t = \hat{y}_t - \sum_k \hat{x}_{k,t} s_{k,t}$$

shows the growth rate of the Solow residual measure of TFP, where a caret denotes log-first-difference, and s is the average of current and lagged factor shares.⁴ Factor shares are computed as revenue shares, where the share of capital is usually computed as 1 minus the shares of the other factors.

Deviations from marginal cost pricing and constant returns will sever the link between the above Solow residual and technological change. An adjustment that was developed by Robert Hall (1990) and applied to microdata by

⁴ For simplicity of exposition, the single output case is shown.

James Levinsohn (1993) and Ann Harrison (1994), for example, uses cost shares as weights (requiring data on the user cost of capital and the assumption of a constant elasticity of demand for output) and explicitly estimates the degree of returns to scale. Other departures between the TFP change index and technological change occur if factor inputs do not receive payments equal to their marginal products or if factors are not utilized fully. Then, allocative and technical inefficiency changes may result in deviations between the Solow residual movements and technological change. A large empirical literature exists that deals with these departures from the basic assumptions and decomposes residuals into efficiency changes and technological change.⁵

A popular method for computing a productivity index for sectoral data is through estimation of a cost function and factor demand equations, a method developed subsequent to the review article by Nadiri (1970) in this journal. The advantage of using this approach to measuring productivity is that the estimated parameters are not biased because of simultaneity of productivity and factor demand. Although the method has sometimes been used for measuring productivity with LMDs (for example, Tybout 1993), the advantages over directly estimating production functions are questionable because identification of the factor demand equations requires variation in factor prices, which are not available at the micro level. Instead, prices vary over time only at the sectoral level.

A highly innovative approach to estimating micro-level productivity while

addressing simultaneity has been developed by Olley and Pakes (1996). Their method, which is being applied frequently in recent LMD research (for example Nina Pavcnik 1998; Levinsohn and Amil Petrin 1999), uses observable micro-level information to correct for the above-mentioned simultaneity bias and to account for self-selection of exiting producers. Their finding—that endogenous exit behavior is important—is corroborated by Bartelsman and Phoebe Dhrymes (1998), who show that capital output elasticities are of the expected magnitude when estimated from an unbalanced panel, whereas they are often insignificant or negative when estimated from a balanced panel. The Olley and Pakes (1996) method hypothesizes that producers observe information about their productivity that is unobservable to the econometrician. Producers use this information along with the level of their capital stock to make decisions on continuation of production and on factor input levels if they decide to continue. By inverting the investment function, one can estimate the unobserved productivity component semi-parametrically as a function of investment. Olley and Pakes show that their resultant estimator for the coefficient on capital is consistent and asymptotically normal.⁶

This subsection has described methodological issues in measuring productivity. The rest of this paper stresses that although some advances have been made (for example, the Olley–Pakes approach) the main contributions from studying LMDs are the insights they afford into the process of economic growth

⁵ For the papers discussed in the present paper, the operational difference between technological change and technical efficiency change, which depends on identifying assumptions in estimation of the production or cost function, does not impinge upon the main results.

⁶ Empirically, the effect of the simultaneity bias on production function parameters is considerable. Levinsohn and Petrin (1999) find that capital output elasticities when controlled for simultaneity are about 0.15 higher on average than when estimated with OLS.

at the micro-level and its relation to more aggregate results.

Limitations of LMDs in Productivity Research. No data set is ideal for every question because of issues related to frequency, coverage, sampling, unit of analysis, time series properties, and missing variables, among others. LMDs are no different. For instance, studies that examine the relationships between R&D and productivity may be appropriately conducted at the firm level, but not at the establishment level. For the LRD, although all the manufacturing establishments of a firm are known in census years, much of the research to date has neglected other parts of the firm, such as headquarters, sales offices, R&D labs, and the other parts that lie outside of manufacturing. Another shortcoming of LMDs is that they tend to focus mainly on manufacturing, a sector whose share of GDP and employment has been decreasing.

Other problems with LMDs arise from productivity measurement, as described above. The difficulties that plague aggregate measures of either total factor productivity or labor productivity—namely, obtaining accurate measures of inputs and outputs—may prove worse at the micro level. In a standard capital, labor, energy, materials (KLEM) model of production, the measurement of each factor has its own issues. Capital stocks or service flows are notoriously difficult to measure at a micro level, given composition effects and the time series of investments required. Fortunately, Baily, Charles Hulten, and David Campbell (1992) find that their results on productivity dispersion and evolution do not change significantly when sophisticated measures of capital are used instead of crude measures based on book value. For labor, the LRD classifies employees only as production or nonproduction workers

rather than providing more detailed information on occupations, education, or experience. Output measurement is difficult as well, with most data being deflated by four-digit industry deflators. Therefore, differentiating between productivity differences and differences in markups is difficult, if not impossible.

A last word of caution: The quality of data collected at the establishment level is largely unknown. Although input and output measures may be noisy at the micro level, aggregate measures may be more precise if the noise is random. Therefore, the question arises of how much of the dispersion in productivity at the micro level is noise and how much is real? The question is briefly addressed in the next section.

3. Some Stylized Facts on Productivity Dispersion and Evolution

Of the basic findings related to productivity and productivity growth uncovered by recent research using micro-data, perhaps most significant is the degree of heterogeneity across establishments and firms in productivity in nearly all industries examined.⁷ The research using LMDs confirms results previously derived from a long history of case studies. For instance, W. Salter (1960) cites evidence of the wide dispersion in labor productivity among plants producing pig-iron: From 1911 through 1926, the best-practice pig-iron facility produced nearly twice as many tons per labor-hour as the industry average. W. Bruce Chew, Kim Clark, and Timothy Bresnahan (1989) analyzed the labor productivity of plants in a multi-plant network, and found intra-firm productivity differences between best and worst plants on the order of 2 to 1, even after controlling for observable

⁷ Unless otherwise noted, productivity refers to TFP.

plant heterogeneity. Broadening the scope to international comparisons, Kent Womack, Daniel Jones, and Daniel Roos (1989) found productivity differences in automobile assembly plants ranging from 13.2 to 78.7 hours of labor per vehicle. Within assembly plants owned and operated by U.S. firms in North America, the authors found productivity ranging from 18.6 to 30.7 hours per vehicle. In a more recent study comparing firms in the United States, Germany, and Japan for a selected group of industries, including automobile assemblers, steel, computers, and beer, McKinsey (1993) found wide gaps in productivity relative to best-practice firms.

Cross-Sectional Dispersion. Several studies have examined the degree of dispersion in productivity using LMDs. In the first such study using the LRD, Dhrymes (1991) constructed measures of TFP—both residuals of production functions and Solow residuals—for a balanced panel of large plants in high-tech industries. The TFP measure for each plant denotes its productivity relative to the average plant in its four-digit industry. The ratio of average TFP for plants in the ninth decile of the productivity distribution relative to the average in the second decile was about 2 to 1 in 1972 and about 2.75 to 1 in 1987.⁸

An interesting question is how much of the measured productivity dispersion reflects real productivity differences across establishments or how much arises from measurement error. Besides statistical errors in reporting nominal values of inputs and outputs, productivity may also be contaminated by qual-

ity differences in output of plants in the same four-digit industry. Further, transitory idiosyncratic productivity shocks would generate measured dispersion even if the permanent component of productivity were the same for all plants. However, evidence on productivity persistence, given below, is not consistent with that hypothesis.

How much of the productivity dispersion is due to dirty data is not known. However, researchers have several reasons to believe that at least a good portion of the dispersion is real. First, as mentioned above, many carefully executed case studies have documented wide productivity dispersion in particular industries. Next, research using LMDs in both developed and developing countries, where statistical errors are likely to vary, have found similar degrees of dispersion. Further, relative productivity between plants has been shown to be correlated with wages, export success, and technology usage—in a sense, the dispersion in productivity appears quite plausible when examined in conjunction with other variables. Finally, high-productivity plants have been found to have higher output growth and are less likely to exit. Having documented the existence of productivity dispersion at a point in time, our next step is to display how the productivity distribution evolves over time.

Productivity Evolution Models. Why is productivity dispersion so great, and what models are useful in thinking about how productivity evolves? One suggestion of Nelson (1981) was that productivity differences reflect the differences in the outcomes of technological bets. Even if plants and firms make the same bets, they may not reap the same rewards. In fact, such bets are fraught with uncertainty. Since Nelson, several models of industry dynamics have formalized these concepts and

⁸ In specific industries, dispersion is also found to be quite large, e.g. in four-digit textile industries, ratios are found between 2 and 4 to 1 (Douglas Dwyer 1996), or in telecommunications equipment, ratios of 3 to 1 (Olley and Pakes 1996). Outside of manufacturing, evidence shows even more dispersion (for example, Nicholas Oulton 1998).

feature firm entry, exit, and productivity evolution. The models vary in their assumptions regarding sunk costs, learning, and the stability of productivity over time. For instance, the Jovanovic (1982) model has firms with a time-invariant efficiency parameter, implying that a firm's productivity will vary initially but then settle down to a constant value. Ericson and Pakes (1995) extend this model so that efficiency is more stochastic: Unlike in the Jovanovic model, negative shocks can cause very productive firms to have losses in efficiency. Lambson's (1991) model takes a different tack and posits that firms choose a technology (based on factor prices), and then become locked into it.

All these theoretical models have some appeal and vary according to assumptions needed to generate productivity dispersion. Very little formal testing has been done as to which of the above models fits the data best. One exception is Pakes and Ericson (1998) who examined the growth patterns of manufacturing and retail establishments in Wisconsin. The authors found that retail firms evolved more like in a Jovanovic (1982) model (once a firm established itself, it was likely to stay), whereas manufacturing firms behaved more in line with the Ericson and Pakes (1995) model (firms are more likely to change in size).

These models can generate quite complex patterns in the evolution of the productivity distribution of firms over time. Statistical models of the productivity evolution may provide a clearer way to describe what is observed in the data. Borrowing from the analysis of Baily, Hulten, and Campbell (1992), four statistical descriptions are distinguished: First, the cross-sectional distribution results from random productivity shocks. Micro-level productivity can be considered a random

perturbance about a common productivity level that has its own time-series evolution. In this case, all firms have the same productivity level except for a transient productivity shock. Industries are comprised of identical firms except for differences (real or measured) due to random shocks. The relative level of productivity of a firm at time t does not help predict productivity at time $t + 1$ —in this case, there would be no persistence. Next, an extension to this model could be where the growth rate in productivity of a firm suffers a random shock relative to the underlying common productivity pattern. In this scenario of random productivity growth, the productivity distribution becomes wider over time (though the specification could be written so that there is an upper limit), and firms that fall below a certain level exit the market. Here the ranking of a firm in the productivity distribution will be correlated with its position in the distribution in the subsequent period. A third statistical description of evolution of the productivity distribution comes from a vintage capital model. In this model, new firms receive the productivity associated with the latest vintage of technology whose evolution is given. Productivity at the new firm subsequently remains unchanged over time, aside from random shocks. Firms exit the industry when their productivity relative to the latest entrants becomes too low. This model has been used frequently in recent theoretical work (Ricardo Caballero and Mohamad Hammour 1994; Jeffrey Campbell 1995) because of its analytical tractability. Finally, the movement of the productivity distribution could be described by a common evolution of heterogeneous plants, with some random level shocks. In this case, the relative differences in productivity between plants are persistent through time.

Productivity Evolution Results—Continuing Establishments. Various parametric and nonparametric methods have been used to describe how plants move over time inside the distribution. One device used in portraying the evolution of productivity is the transition matrix, which shows the probability that a plant that resided in a certain productivity percentile in an earlier period moves to other percentiles in a later period. In Bartelsman and Dhrymes (1998), the year-to-year transition probabilities show a very high degree of persistence; after five years, more than one-third of the plants remains in the same quintile. Baily, Hulten, and Campbell (1992) used a sample of all plants in 23 four-digit SIC industries to describe the dynamics of plant productivity.⁹ In this paper, the transition probabilities are weighted by employment in each plant and thus show the percentage of employment in plants moving between productivity quintiles. What becomes clear from the transition measures is that about 20 percent of employment remains in its original position after ten years, that the probability of death decreases with productivity rank, and that entrants are not predominantly in the top quintile at the later date. Further tests of persistence compare the rank statistics between periods using the Wilcoxon test; in the Baily, Hulten, and Campbell study, plants initially in the top two quintiles were still ranked far above average ten years later.

Productivity Evolution Results—Entry and Exit. The dispersion of productivity is wide for a balanced panel of

continuously operating, large plants. Once small plants are allowed in the sample under review, as well as plants that exit or enter within the period, the variance of the productivity distribution is likely to increase. Making precise estimates of the effect is difficult because obtaining capital stock estimates for small plants and for plants that remain in the sample only for short periods is very difficult. Nonetheless, research seems to show that the productivity of entering plants is below that of the average incumbent.¹⁰ A few years after entry, a disproportionate number of entrants have moved both to the highest and within the lowest productivity quintiles. This pattern of productivity and output is explained nicely by John Baldwin (1995) in his study of the Canadian manufacturing sector. Entrants are much smaller than the average incumbent, and about half die within the first decade. If they survive, entrants reach average productivity in about a decade, although they are still smaller than the average. The probability that a given plant will die turns out to be inversely related to both productivity and plant size. Essentially, the pattern is one of weeding out the unsuccessful entrants and of nurturing (both in terms of output and productivity) the successful ones. The significance of the contribution of the weeding out and nurturing of entrants, as well as the beneficial contribution of exit by low productivity plants to aggregate productivity growth, will be discussed below.

Productivity Evolution Results—Cyclicalities. The question of whether productivity moves procyclically—and, if so, why—is important for understanding the source of business cycles. Real business cycle proponents contend that

⁹ As referenced in Baily, Hulten, and Campbell (1992), Bartelsman and Dhrymes (1998) use similar methods and reach similar conclusions concerning productivity dynamics. The study by Baily, Hulten, and Campbell is more comprehensive and will be used to describe the common findings.

¹⁰ This finding is at odds with the vintage capital story of productivity dispersion.

business cycles are caused by intertemporal shifts in production induced by positive productivity shocks. Alternatively, movements over time in measured productivity could be caused by changes in output. Explanations for procyclical Solow residuals range from increasing returns to scale, varying markups (Hall 1990), and external economies (Bartelsman, Caballero, and Richard Lyons 1994) to cyclical utilization (Susanto Basu and John Fernald 1997).

Recent evidence from microdata reveals that aggregate labor productivity is less procyclical than within-plant productivity, because of reallocation of market share from low-productivity to higher-productivity plants during downswings (Baily, Bartelsman, and John Haltiwanger 1999). At the micro level, researchers have generally found constant returns to scale, while evidence on time-varying markups is difficult to interpret because of lacking micro-level price data.

Microlevel data provide a nice setting to evaluate the role of factor utilization in generating procyclical productivity. Firms that expect to be downsizing over time would not let workers sit idle during a cyclical downturn but would make use of the opportunity to dismiss them. However, the same firms would certainly attempt to utilize workers harder during a temporary upswing rather than add extra workers. Similarly, firms that expect to grow over time would hire during an unexpected temporary boom but likely would let workers sit idle during a temporary downturn. Recent evidence is at odds with this story of changing factor utilization as the main source of procyclicality. In U.S. manufacturing, long-run downsizers exhibit more procyclicality than upsizers. Procyclicality of long-run upsizing plants during cyclical downturns is not stronger than during upturns, while

procyclicality of long-run downsizing plants during cyclical upturns is no greater than during downturns (Baily, Bartelsman, and Haltiwanger 1999).

Overall, the procyclical productivity puzzle is not yet solved. The asymmetry in measured procyclicality between upsizers and downsizers could point toward asymmetric adjustment costs. Alternatively, the types of shocks that hit producers may affect long-run outcomes at the microlevel, leaving the door open for a link between technology shocks and output.

Productivity Evolution Results—Summary. After all is said and done, which of the statistical models presented by Baily, Hulten, and Campbell (1992) best represents the patterns of productivity evolution? The random productivity growth pattern, with the widening distribution of productivity over time, is the least representative. Productivity as a random shock about a common evolution has some support in the data. There does appear to be reversion to the mean, although it may result from noisy data. The pattern of evolving common productivity among heterogeneous plants is also supported to some extent because the amount of persistence in the location of a plant in the distribution is large. The vintage capital model fits with the fact that exit probabilities increase with age and decrease with productivity. However, a strict vintage story is contradicted by the fact that new plants do not enter at the top of the productivity distribution.¹¹ All told, no single statistical model adequately captures all the variation in productivity of production units in any industry.

¹¹ Another way to think of vintage is the age of capital in a plant. Laura Power (1998) examined productivity growth in manufacturing establishments after periods of large investments. She found that productivity growth did not increase faster than average after a large amount of new capital was installed.

Aggregate versus Average Productivity Growth. A primary reason for such interest in examining micro productivity growth is to gain a better understanding of aggregate productivity growth. Aggregate productivity growth is a weighted average of plant-level productivity growth, where the weights are related to the importance of the plant in the industry:

$$\Delta \Pi_t = \sum_i s_{it} \Delta \pi_{it}$$

where Δ denotes first difference, π is (log) productivity, and s is the share of firm i (or establishment i) in the total. Different authors have used various methods for computing the contributions of compositional shifts of output and within-plant productivity growth to aggregate productivity, but the results point in the same direction: Aggregate TFP growth in the periods studied is significantly affected by compositional changes, such as shifts in output toward plants with relatively high TFP.

In a recent paper, Lucia Foster, Haltiwanger, and C. J. Krizan (1998) review the computations used in empirical studies that decompose aggregate TFP growth into components related to within-plant TFP growth, reallocation, and the effects of exit and entry. Their preferred decomposition is a modified version of that used by Baily, Bartelsman, and Haltiwanger (1996) and provides a framework to interpret the seemingly disparate findings in the literature:

$$\begin{aligned} \Delta \Pi_t = & \sum_{i \in C} s_{it-1} \Delta \pi_{it} \\ & + \sum_{i \in C} (\pi_{it} - \Pi_{t-1}) \Delta s_{it} + \sum_{i \in C} \Delta \pi_{it} \Delta s_{it} \\ & + \sum_{i \in N} s_{it} (\pi_{it} - \Pi_{t-1}) \\ & - \sum_{i \in X} s_{it-1} (\pi_{it} - \Pi_{t-1}) \end{aligned}$$

In the decomposition, aggregate productivity growth between two periods is

composed of five parts. The first three add up to the contribution of continuing plants, and the last two, of entry and exit, respectively. The labor shares, s , are computed over all plants in the relevant period. The five components distinguished are (1) a within-plant effect—within-plant growth weighted by initial output shares; (2) a between-plant effect—changing output shares weighted by the deviation of initial plant (log) TFP and initial (log) industry TFP; (3) a covariance term—a sum of plant TFP growth times plant share change; (4) an entry effect—a year-end share-weighted sum of the difference between (log) TFP of entering plants and initial (log) industry TFP; and (5) an exit effect—an initial-share-weighted sum of the difference between initial (log) TFP of exiting plants and initial (log) industry TFP.

The between-plant and the entry and exit terms use the deviation between plant productivity and the industry average in the initial period. A continuously operating firm with an increasing share makes a positive contribution to aggregate productivity only if it initially has higher productivity than the industry average. Entering (exiting) plants contribute only if they have lower (higher) productivity than the initial average. This treatment of births and deaths ensures that the contribution to the aggregate does not arise because the entering plants are larger than exiting plants, but because of productivity differences.

Table 1 shows that over a ten-year period, one-quarter of TFP growth is the result of net entry and exit. The within-plant component is almost half.¹²

¹² Bartelsman and Dhrymes (1998) find that only a fraction of growth is from the within-plant component. However, their decomposition contrasts unweighted average productivity growth with the aggregate. In general, larger plants are more productive, a finding that is captured in the within-plant component in the table above because plants are weighted by initial share.

TABLE 1
DECOMPOSITION OF TFP GROWTH FOR U.S. MANUFACTURING ESTABLISHMENTS, SELECTED PERIODS

Census period	Total growth	Within-plant share	Between-plant share	Cross-plant share	Net entry share
1977–87	10.24	0.48	–0.08	0.34	0.26
1977–82	2.70	–.09	–0.33	1.16	0.25
1982–87	7.32	0.52	–0.18	0.51	0.14

Notes: Tabulations from LRD by Foster, Haltiwanger, and Krizan (1998).

The covariance term captures the dynamic interaction among continuing plants, with the contribution to this term being positive when productivity growth and changes in share move in the same direction. Decompositions such as this can be modified to split the data in other interesting ways. For example, Baldwin (1995) wondered about the contribution of plants with declining market shares to aggregate productivity growth. The (weighted) sums in the decomposition can be split into a summation over plants with long-run increasing and decreasing market share. Baldwin considered the productivity advances made by the least successful plants to be exogenous and found that such disembodied advances constitute only about 4 percent of aggregate growth.

A revealing decomposition of labor productivity is given by Baily, Bartelsman, and Haltiwanger (1996). Although at the aggregate manufacturing level rising labor productivity is accompanied by reductions in labor input, at the micro-level this correlation does not hold. Among plants with increasing labor productivity between 1977 and 1987, equal proportions increased employment as decreased employment. Further, the proportion of downsizing plants that experienced productivity increases was no different than the proportion of growing plants that experienced rising productivity. Overall, about a third of aggregate labor productivity growth of 3

percent per year is attributable to up-sizers, about a third to downsizers, and a third to the net effect of entry and exit.¹³

Stylized Facts and Fallacies. Researchers using LMDs have documented that productivity levels are quite dispersed, that productivity differences between plants may be very persistent, that entry and exit of plants with different productivity levels is an important source of productivity growth, and that plants' long-run employment changes and productivity changes are not correlated. The existence of productivity heterogeneity, even among producers of comparable products with comparable equipment, has forced analysts to rethink and reassess some old "truths" that find no support in the microdata.

For instance, these results begin to cast doubt on the appropriateness of an aggregate production function that is based on a representative firm. Industry output is not produced with industry inputs in such an orderly fashion. Factor output elasticity, which is considered a structural property of a production function, does not represent a property of "industry technology." If, for example, labor input in an industry were to double, the resulting change in industry

¹³ Similar calculations have been made for manufacturing in various countries, with similar results (Lutz Bellmann and Tito Boeri 1998; Bartelsman, George van Leeuwen, and Henry Nieuwenhuijsen 1998).

output would depend not only on the production technology at each plant but also on the allocation of inputs and output among plants within the industry before and after the change. The properties associated with an aggregate production function, such as factor substitution and scale elasticities, do not represent marginal responses of the industry (nor of a plant). Instead, the elasticities capture both some average (over time and across plants) of micro-level production technology and the effects of past changes in composition within the industry.

A clear example of the problems of interpretation of substitution elasticities can be found in Doms (1993). Using linked data from the Manufacturing Energy Consumption database and the LRD, he traced the response of an industry's energy inputs to a change in a certain fuel's price. Changing fuel prices can alter plant-level decisions of incurring a fixed cost to purchase fuel-using equipment of a certain type. However, given choices of fuel-use equipment made in the past, which is shown to be very heterogeneous, plants do not have much scope for substitution. Yet the industry may exhibit considerable substitution because of shifts in industry output from plants with relatively high use of the high-cost fuel towards plants that use relatively less.

Another example of mistaken conclusions that can be drawn from analysis at an aggregate level is pointed out by Tybout and Daniel Westbrook (1996). Although sectoral and macroeconomic estimates often display increasing returns to scale, Tybout and Westbrook show that scale effects are an unlikely source of welfare gains following trade liberalization. Using Mexican plant-level manufacturing data, they showed that scale economies are present mostly for small plants and that larger plants gen-

erally exhibit constant returns. Because changes at large plants account for most fluctuations in industry output, scale economies cannot be exploited at the margin of interest.

An implication of the above-mentioned problems of elasticities estimated with aggregate data is that the elasticities may be time varying and nonlinear and that they do not reflect structural characteristics of the production technology in any case. A lesson to be learned is that one cannot rely on aggregate elasticities in order to compute the marginal response to changes in relevant variables. Rather, one must understand the micro-level technology and the mechanism of compositional response to changes in the variable to calculate marginal responses. Of course, this puts an inordinate burden on a researcher trying to assess, for example, the impact on information technology investment of changes in the relative cost of capital. Advances, however, are being made in micro-macro models that do not require complete microdata, but rely instead on information contained in time series of higher moments of the cross-sectional distribution of relevant variables (see Haltiwanger 1997, for a review of work that has examined the use of microdata in analyzing aggregate fluctuations.)

4. Factors behind the Patterns

What makes some plants more productive than others? What are some of the forces that influence the patterns described in the previous section? Some of the factors that are thought to be important include ownership, quality of the workforce, technology, international exposure, and the regulatory environment. Much of the research that examines the relationship between some factor (say, ownership or human capital)

and productivity growth is limited to showing a correlation, and does not determine causality. Showing causality is difficult in this line of research, as it is in most empirical work. Nonetheless, studies using such microdata arguably have improved our understanding of the effects of (de)regulation and competition, the story behind management/ownership changes, the importance of technology and human capital, and the stimulus provided by international exposure.

Although these factors all are important and significant, unfortunately they do not explain a significant portion of the heterogeneity documented in the previous section. Perhaps this is not surprising, as pointed out by Griliches and Jacques Mairesse (1983): "The simple production function model, even when augmented by additional variables and further nonlinear terms, is at best just an approximation to a much more complex and changing reality at the firm, product, and factory floor level." At the micro level, productivity remains very much a measure of our ignorance.

Regulation. Policy evaluation is particularly difficult to undertake at the micro level. Regulatory policy may affect not only the decisions firms make today but also the market structure tomorrow, by altering the incentives for innovating, investing, market entry, and gaining market share. To evaluate a given policy on an industry, one would need to estimate a model with endogenous industry structure, because a firm's actions depend on, and affect, industry structure. So far, such models have been simulated, for example by Pakes and Paul McGuire (1994) and Hugo Hopenhayn and Richard Rogerson (1993), but no attempts have been made to estimate econometrically the effects of policy changes on industry structure.

The LRD has been used to examine a

market before and after deregulation: Olley and Pakes (1996) estimate firms' exit and investment behavior in the telecommunications equipment industry during successive stages of deregulation in the 1960s, 1970s, and 1980s and describe changes in industry structure. The basic finding is that considerable resource reallocation coincided with deregulation. For instance, in the five years after the 1968 easing of regulations for interconnection of private equipment to the public network, the number of plants and firms in the industry almost doubled. The amount of reallocation, including that which resulted from exit, increased dramatically after the AT&T divestiture; many of the plants that manufactured telecommunications equipment in 1972 were closed by 1987. Since 1975, most of the productivity growth in the industry has arisen from resource reallocation, which has been boosted by particularly high exit probabilities for low productivity plants.

In summary, the telecommunications equipment industry provided a good case study on the way changes in regulation contributed to productivity growth. First, entry barriers broke down. The breakdown altered choices of producers and potential producers regarding their innovative activity, input choices, and production volumes. Many new plants and firms entered the industry, while inefficient plants decided to stop production. The interactions in the marketplace became different over time as vertical production chains were severed and competitors entered the market. Finally, aggregate productivity grew through the changing market shares and interactions with productivity of the enlarged field of competitors.

Management/Ownership. How important is managerial ability in understanding productivity differences? On a

theoretical level, Jovanovic's model can be thought of as a manager skill-model—better managers possess higher efficiency parameters and have higher productivity. In contrast, in Lucas's (1978) model, which has an equilibrium where better managers control larger enterprises, labor productivity is the same across enterprises because of the diminishing returns to skill. On a practical level, managerial quality may be an important factor behind productivity heterogeneity. Further, leveraging of managerial quality through better matching between managers and production operations may have been the crucial motive underlying the wave of mergers and acquisitions in the 1980s, rather than the less-efficient motive of empire building. Testing the role of managerial quality is problematic because data on manager quality cannot be directly collected.¹⁴ Nonetheless, a series of results from papers using the LRD yield some important insights.

It has been found that productivity at an establishment is positively related to the productivity of the firm to which it belongs, both in levels and growth rates. Baily, Hulten, and Campbell (1992) state, "Well-run firms will be able to transfer those skills to their plants by training managers, giving advice, and transferring technology, good product design, and production methods." Just how much of the variance in the cross-section of productivity or productivity growth can be attributable to ownership? Although "firm effects" are thought to be important by researchers (many regressions control for the firm to which an establishment belongs), the importance of firm effects has yet to be adequately quantified.

¹⁴ Differences in managerial ability may result in differences in technology usage and in the quality of workers, both of which are discussed separately below.

Perhaps the area in which ownership has received the most attention is the effect of mergers. After a wave of large conglomerate mergers in the 1960s, researchers speculated that changes in ownership may have had detrimental effects on productivity, as the motivation for the mergers may not have been profit maximization but empire building. The issue was revisited in the 1970s and 1980s with the onset of a new wave of mergers. On the other hand, some researchers argued that changes in ownership may increase efficiency if, by changing ownership, better matches were formed between enterprises and management. To test which story predominated, several studies used the LRD to examine the productivity of establishments before and after changes in ownership (see, for example, Frank Lichtenberg 1992; McGuckin and Sang Nguyen 1995).

Both sets of studies found that establishments enjoyed above-average productivity growth for several years following a change in ownership. Lichtenberg and Donald Siegel (1990a) extended their analysis by examining productivity changes by the type of ownership change—leveraged buyouts (LBOs), management buyouts (MBOs), or "other" changes—hypothesizing that the different types of mergers may have had different motivating factors. They found that establishments that underwent LBOs and MBOs were more productive than average before the buyout and that productivity of these establishments increased faster after the buyout than the productivity of establishments that underwent other types of changes in ownership.

One reason that mergers may increase productivity is that they reduce corporate overhead. To test this hypothesis, Lichtenberg and Siegel (1990b) also examined the relationship between

ownership changes and changes in central office personnel, including R&D personnel. They found that ownership changes were related to a reduction in the share of employment at auxiliary offices. However, they did not find a statistically significant relationship between ownership change and R&D personnel.

These results had a significant effect on the debate concerning the effects of mergers on productivity and the motivation behind mergers. The studies using the LRD introduced results from large samples into an area of study where anecdotal evidence, or results from small and unrepresentative samples, had been the norm.¹⁵ An additional advantage of using the LRD in these studies was that it allowed examination of the parts of firms that changed ownership, which is important given the large degree of partial divestitures and acquisitions. However, as in any area of empirical work, there are caveats. A fundamental difficulty in comparing two groups of establishments is selection: the establishments that undergo ownership change are not a random sample from the population. Hence, making statements about causality becomes problematic. McGuckin and Nguyen examined in detail how those establishments that undergo ownership change differed from the population; in the absence of a randomized experiment, they could never be certain whether changes in ownership caused productivity to increase, although the evidence presented dovetailed nicely with the matching theory.

In summary, the story line behind ownership/management and productivity may be one close in spirit to Nelson, in which market participants grope to find

more efficient means of meeting market demand. Firms' choices of technology, inputs, and production are made by management. Different management may be the means toward better choices. An unexplored area for future research is the effect that mergers may have on market interactions through increases in market concentration.

Technology and Human Capital. Another reason for wide productivity differences across establishments and firms is heterogeneity in the technology—specifically, the physical and human capital—used to produce goods. At a point in time, not all firms producing similar goods employ the same capital: Some firms will use the latest and best technology, while others use more traditional methods. Nelson (1981) emphasized the importance of technological change in firm productivity growth and of understanding the way technology is generated and distributed through firms. Plenty of studies have documented the correlation between some measure of technology and productivity at the micro level. (For R&D and productivity: Lichtenberg and Siegel 1991; Bronwyn Hall and Mairesse 1995. For advanced technology and productivity: Dunne 1994; McGuckin, Mary Streitwieser, and Doms 1998; Lucy Chenells and John Van Reenen 1998. For computers and productivity: Lichtenberg 1995; Erik Brynjolfsson and Lorin Hitt 1995.) Unfortunately, these findings may not tell us much about why firms or plants with high productivity adopt the technology, or vice-versa.

A first step in examining how technology affects productivity is to examine how it diffuses through the economy. Many case studies have found an S-shaped pattern of diffusion: relatively few early adopters are followed by greater and greater use by the remaining population. Dunne (1994) used a

¹⁵ An advantage of the LRD over data sets such as Compustat is that the LRD contains data on publicly and privately owned firms whereas Compustat has data only for publicly traded companies.

supplemental survey on technology usage at manufacturing establishments and found great heterogeneity even after controlling for a large list of factors, such as industry, size, region, age, and so forth.¹⁶ After controlling for plant size, technology usage was found not to be correlated with age. This finding is interesting because, as mentioned previously, several models are based on the assumption that new technology is introduced through new plants.

Because the technology data in the Dunne study was from a cross-sectional survey, questions of causality were difficult to address. For instance, did plants become larger after adopting advanced technologies, or were large plants (which may have acted as a proxy for productivity and successfulness) better able to adopt new technologies? Bartelsman, van Leeuwen, and Nieuwenhuijsen (1998) found that technology-intensive manufacturing firms in the Netherlands were more likely to increase productivity and add workers, even after controlling for the technology-adoption decision. However, identification of the model is based on functional form and temporal restrictions of the adoption model rather than on exogenous instruments.

Examining the relationship between productivity and advanced technology by itself can lead to only limited insights. Technology usage is related to other variables, such as labor quality. In fact, there has been much speculation that technological change is skill-biased. Detailed data on workers exist for a small subset of establishments in the LRD.¹⁷ Using these data in conjunction

with the advanced technology usage data, Doms, Dunne, and Troske (1997) investigated the relation between human capital and technology. They found a strong, positive relationship between the number of advanced technologies used and the skills of both production and non-production workers. Perhaps more interesting, they found that plants that used advanced technologies in 1988 also had above-average productivity in 1972. How could this be? If technology and worker skills are complements, then firms that have above-average worker skills may be better able to adopt the latest technologies, reinforcing their lead over rivals. This explanation is broadly consistent with the persistence in productivity that was discussed earlier.¹⁸

In summary, the issue of technology is a component of the many choices facing individual producers. By trying to isolate the effects of technology choice on productivity, one may obscure the rich set of concerns facing producers as they attempt to meet the market. It is unlikely that the high measured rates of return to IT investment (for example, Brynjolfsson and Hitt 1995) or R&D (for example, Hall and Mairesse 1995) or advanced technology usage (Doms, Dunne, and Troske 1997) point to any unexploited arbitrage opportunities at the micro level. More likely, omitted variables, such as complementary investments in organization capital, explain the difference between measured and market returns (as explored in theoretical work by Michael Kremer and Eric Maskin 1996).

International Exposure. One of the

¹⁶ Dunne's study was based on a statistically drawn sample of over 8,000 establishments from a wide variety of industries.

¹⁷ A subset of workers from the 1990 Census of Population has been linked to the establishments where they work. Such linked worker-firm data sets are available for larger portions of the popula-

tion or total employment in France, Denmark, Finland, and Norway. Workers can even be tracked over time as they move from firm to firm.

¹⁸ Similar findings to Doms, Dunne, and Troske (1997) was found by Horst Entorf and Francis Kramarz (1998) using French data. They found that above average workers were selected to use new technologies.

more interesting issues is understanding the way market interactions affect productivity. Recently, research has explored how exposure to foreign markets relates to producers' choices, productivity dispersion, and market share. Plant-level exporting has received a good deal of attention, driven in part by the strong relationship found between exporting and productivity (see Andrew Bernard and Jensen 1995). Exporting may be related to productivity and productivity growth in several ways. First, firms that have high productivity may hold comparative advantage and are more likely to overcome the obstacles to exporting than less-efficient firms. Second, firms that export increase their exposure to other firms and may learn about available production technologies. Third, exporting may allow firms to operate at a more-efficient scale, moving down their average cost curves.

Many studies have found a positive relationship between exporting and productivity in a cross-section. These results, in addition to similar results found for wages and exporting, led some to advocate that governments take a more active role in promoting exporting (see Sofronis Clerides, Saul Lach, and Tybout 1998). Going a step further, Bernard and Jensen (1999) addressed the question of whether exporters increase their productivity after entering foreign markets and whether firms that hit upon successful strategies (that is, firms that increase their productivity) are more likely to export. They found that relatively productive firms are likely to export, but they found little change in efficiency after exporting had begun. These results are similar to those of others who examined microdata for three semi-industrialized countries.¹⁹

Another aspect of international competition that has been examined is foreign direct investment. According to Doms and Jensen (1998), foreign-owned manufacturing establishments have 2.3 to 3.7 percent higher total factor productivity than domestically owned establishments, even after controlling for a host of variables, such as detailed industry, size, age, and state. This result is consistent with the theory that foreign firms that establish operations in the United States have advantages in efficiency and product design. The analysis then goes a step further by examining not only foreign multinationals but also U.S. firms that have overseas assets (U.S. multinationals). As a group, the domestic production units of U.S. multinationals are the most productive plants. However, foreign-owned plants are more productive than U.S.-owned plants with no overseas assets.

In summary, the papers discussed above seem to point out the way market interactions stimulate better choices by producers at the micro level. The ability to engage in activities in different markets allows multinational firms to improve the match between a production unit and the market it operates in. Exposure to foreign markets is facilitated by more efficient operations. The link between exposure to foreign markets and better micro-level decisions has not yet been established.

5. *Future Directions*

This paper has reviewed some of the productivity research that has been conducted using LMDs. That research has evolved from the documentation of micro-level productivity growth, to the examination of the factors associated with

¹⁹ Tybout and Westbrook (1996) examine the issue of trade liberalization and increasing effi-

ciency due to exploiting plant-level returns to scale and find that there is not much of a scale-economy effect.

productivity, and finally to more rigorous analyses that have begun to explore causal relationships. The results have been new insights into the important, outstanding questions in productivity analysis, such as what the sources behind productivity growth are and what effect institutions have on productivity growth. Perhaps the most significant contribution LMDs have made is to revisit the ideas of heterogeneity and Schumpeter's creative destruction. The premise underlying much of this work is that tremendous differences in productivity exist across firms and that understanding the reasons for these large disparities may yield insights into the way productivity growth occurs. The argument echoes Nelson's (1981) earlier analysis, but it now can be addressed better quantitatively.

There are several areas in productivity research where LMDs have not yet been used and others in which LMDs can be used more effectively. These are areas for future work.

1. **Reasons for heterogeneity:** The work reviewed documents much of the anecdotal evidence on firm heterogeneity that Nelson (1981) also described. Emerging in this line of research is the high correlation among technology, human capital, management, and other factors. Research that uses data sets with only one or two of these additional variables may simply be picking up the role of some omitted variable. Work to address causal relationships between productivity and a host of factors has only just begun and needs to progress further. The complexity of the interactions between the factors correlated with productivity growth, such as worker skills, technology usage, management techniques, and quality of output, is calling out for some elegant theoretical explanations.
2. **Beyond manufacturing:** A disproportionate amount of research using LMDs has focused on the manufacturing sector. The lack of attention to nonmanufacturing arises mostly because of difficulties in defining output and measures of inputs. In the near term, more work could be done to include other parts of the manufacturing firm in studies that rely on manufacturing establishment data. For instance, data on auxiliary offices or sales offices of manufacturing firms could be included. As the share of employment in manufacturing continues to decrease, the need for understanding productivity outside of manufacturing will become even more imperative. A step in this direction is the study by Foster, Haltiwanger, and Krizan (1998) which uses census data to examine productivity dynamics in the automobile repair industry. At the Center for Economic Studies at the Census Bureau, work is now being done to develop LRD-type data sets for nonmanufacturing industries.
3. **Another direction to pursue on the data front** is to link data on workers to the establishments in which they work, given the strong relationships found between technology use and worker skills. Doms, Dunne, and Troske (1997) is one example of this work in the U.S., and Entorf and Kramarz (1998) is an example using French data.
4. **Data quality:** Because LMDs provide a large number of observations, and hence lower standard errors, much of the research using LMDs has not explored data quality. This neglect has left open the questions of how much of the heterogeneity is real and how much comes from errors in variables; nor has much attention been given to the statistical properties of linked data sets.

5. Price taking and market interactions: Nearly all studies using LMDs have been conducted in price-taking frameworks: that is, the decisions and behavior of firms or establishments have been examined under the assumption of no interactions among firms and no interactions between establishments within the same firm. The decision to use certain technologies or to increase output depends on the actions of competing firms. Modeling these interactions would be difficult, and Steven Berry, Levinsohn, and Pakes (1995) have provided an example of how complex modeling market equilibrium can become. Nonetheless, the extension to a more-general equilibrium framework seems to be a logical next step.
6. Cross-country comparisons: Many of the stylized facts described in Section 3, such as persistence in productivity dispersion, are common features of LMDs across countries. However, exact magnitudes of the statistical features vary. The cross-country variance in measures built up from the microdata can be used to help identify the role played by economic policy and institutions, which vary widely across countries. This type of research will be difficult to undertake because one needs to control the data for differences in survey methodology. Further, because the data need to be analyzed separately at the national sites that allow access, methods need to be developed for "pooling" relevant features of the data without violating disclosure rules.
7. Increased micro-macro linkages: Greater attention should be paid to the aggregate implications of the findings from microdata and to micro-implications of findings at the aggregate level. Several microdata studies do this admirably, while others simply

report their findings without any reference to how important those findings are to aggregate productivity growth. For instance, LMDs have not been used to specifically address the long-standing issue of the productivity slowdown that started in the 1970s. And, looking ahead, LMD researchers will want to be positioned to contribute to the emerging debate over whether or not aggregate productivity accelerated in the 1990s.

REFERENCES

- Baily, Martin N. 1986. "Productivity Growth and Materials Use in U.S. Manufacturing," *Quart. J. Econ.* 101:1, pp. 185-95.
- Baily, Martin N.; Eric J. Bartelsman, and John C. Haltiwanger. 1996. "Downsizing and Productivity Growth: Myth or Reality?" *Small Bus. Econ.* 8:4, pp. 259-78.
- . 1999. "Labor Productivity: Structural Change and Cyclical Dynamics," mimeo, U. Maryland.
- Baily, Martin N.; Charles Hulten, and David Campbell. 1992. "The Distribution of Productivity in Manufacturing Plants," *Brookings Pap. Econ. Act.: Microeconomics*, Washington, DC, pp. 187-249.
- Baldwin, John R. 1995. *The Dynamics of Industrial Competition: A North American Perspective*. NY: Cambridge U. Press.
- Balk, Bert M. 1998. "Industrial Price, Quantity, and Productivity Indices: The Micro-Economic Theory and an Application," Boston; Dordrecht, London: Kluwer Academic.
- Bartelsman, Eric J.; Ricardo J. Caballero, and Richard K. Lyons. 1994. "Customer- and Supplier-Driven Externalities," *Amer. Econ. Rev.* 84:4, pp. 1075-84.
- Bartelsman, Eric J. and Phoebus J. Dhrymes. 1998. "Productivity Dynamics: U.S. Manufacturing Plants 1972-1986," *J. Productiv. Anal.* 9:1 pp. 5-34.
- Bartelsman, Eric J. and Wayne Gray. 1996. "The NBER Manufacturing Productivity Database," NBER Technical Working Paper 205.
- Bartelsman, Eric J.; George van Leeuwen, and Henry Nieuwenhuijsen. 1998. "Adoption of Advanced Manufacturing Technology and Firm Performance in the Netherlands," *Econ. Innovation New Tech.* 6:4, pp. 291-312.
- Basu, Susanto and John Fernald. 1997. "Returns to Scale in U.S. Production: Estimates and Implications," *J. Polit. Econ.* 105:2, pp. 249-83.
- Bellmann, Lutz and Tito Boeri. 1998. "Internal and External Adjustment to New Technologies," *Econ. Innovation New Tech.* 5:2-4, pp. 325-43.

- Bernard, Andrew and J. Bradford Jensen. 1995. "Exporters, Jobs, and Wages in U.S. Manufacturing: 1976-1987," *Brookings Pap. Econ. Act.: Microeconomics*, Washington, DC, pp. 67-112.
- . 1999. "Exceptional Exporter Performance: Cause, Effects, or Both?" *J. Int. Econ.* 47:1, pp. 1-25.
- Berry, Steven; James Levinsohn, and Ariel Pakes. 1995. "Automobile Prices in Market Equilibrium," *Econometrica*, 63:4, pp. 841-90.
- Brynjolfsson, Erik and Lorin Hitt. 1995. "Information Technology as a Factor of Production: The Role of Differences among Firms," *Econ. Innovation New Tech.* 3:3-4, pp. 183-99.
- Caballero, Ricardo J. and Mohamad Hammour. 1994. "The Cleansing Effect of Recessions," *Amer. Econ. Rev.* 84:5, pp. 1350-68.
- Campbell, Jeffrey R. 1997. "Entry, Exit, Technology and Business Cycles," NBER working paper W5955.
- Caves, Richard E. 1998. "Industrial Organization and New Findings on the Turnover and Mobility of Firms," *J. Econ. Lit.* 36:4, p.p. 1947-82.
- Chennells, Lucy and John Van Reenen. 1998. "Establishment Level Earnings, Technology and the Growth of Inequality: Evidence from Britain," *Econ. Innovation New Tech.* 5:2-4, pp. 139-64.
- Chew, W. Bruce; Kim B. Clark, and Timothy Bresnahan. 1989. "Measurement, Coordination, and Learning in a Multiplant Network," Stanford Center Econ. Policy Research working paper 174.
- Clerides, Sofronis; Saul Lach, and James R. Tybout. 1998. "Is 'Learning by Exporting' Important? Micro Dynamic Evidence from Colombia, Mexico and Morocco," *Quart. J. Econ.* 113:3, pp. 903-47.
- Doms, Mark E. 1993. "Inter Fuel Substitution and Energy Technology Heterogeneity in U.S. Manufacturing," U.S. Bureau of Census, Center Econ. Studies discussion paper 93-5.
- Doms, Mark E.; Timothy Dunne, and Kenneth Troske. 1997. "Workers, Wages and Technology," *Quart. J. Econ.* 112:1, pp. 253-90.
- Doms, Mark E. and J. Bradford Jensen. 1998. "Comparing Wages, Skills, and Productivity Between Domestic and Foreign Owned Manufacturing Establishments in the United States," in *Geography and Ownership as Bases for Economic Accounting*. R. Baldwin; R. Lipsey, and J. Richardson, eds. Chicago: U. Chicago Press.
- Dhrymes, Phoebus. 1991. "The Structure of Production Technology: Productivity and Aggregation Effects," Columbia U. Econ. Dept. working paper 551.
- Dunne, Timothy. 1994. "Patterns of Technology Usage in U.S. Manufacturing Plants," *RAND J. Econ.* 25:3, pp. 488-99.
- Dwyer, Douglas. 1996. *The Evolution of an Industry: Theory Applied to Plant-Level Data*, Ph.D. dissertation, Columbia U.
- Entorf, Horst and Francis Kramarz. 1998. "New Technologies, Wages, and Worker Selection," *Econ. Innovation New Tech.* 5:2-4, pp. 165-97.
- Ericson, Richard and Ariel Pakes. 1995. "Markov Perfect Industry Dynamics: A Framework for Empirical Work," *Rev. Econ. Stud.* 62:1, pp. 53-82.
- Foster, Lucia; John C. Haltiwanger, and C. J. Krizan. 1998. "Aggregate Productivity Growth: Lessons from Microeconomic Evidence," NBER working paper W6803.
- Griliches, Zvi and Jacques Mairesse. 1983. "Comparing Productivity Growth: An Exploration of French and U.S. Industrial and Firm Data," *Europ. Econ. Rev.* 21:1-2, pp. 89-119.
- Griliches, Zvi and Vidar Ringstad. 1971. *Economics of Scale and the Form of the Production Function*. Amsterdam: North Holland.
- Hall, Bronwyn and Jacques Mairesse. 1995. "Exploring the Relationship between R&D and Productivity in French Manufacturing Firms," *J. Econometrics*, 65:1, pp. 263-93.
- Hall, Robert E. 1990. "Invariance Properties of Solow's Productivity Residual," in *Growth/Productivity/Unemployment*. P. Diamond, ed. Cambridge: MIT Press.
- Haltiwanger, John C. 1997. "Measuring and Analyzing Aggregate Fluctuations: The Importance of Building from Microeconomic Evidence," *Fed. Res. Bank St. Louis Rev.* 79:3, pp. 55-77.
- Harrison, Ann. 1994. "Productivity, Imperfect Competition and Trade Reform: Theory and Evidence," *J. Int. Econ.* 36:1-2, pp. 53-73.
- Hopenhayn, Hugo A. and Richard Rogerson. 1993. "Job Turnover and Policy Evaluation: A General Equilibrium Analysis," *J. Polit. Econ.* 101:5, pp. 915-38.
- Jovanovic, Boyan. 1982. "Selection and Evolution of Industry," *Econometrica*, 50:3, pp. 25-43.
- Klette, Tor J. and Zvi Griliches. 1996. "The Inconsistency of Common Scale Estimators When Output Prices Are Unobserved and Endogenous," *J. Appl. Econometrics*, 11:4, pp. 343-61.
- Kremer, Michael and Eric Maskin. 1996. "Wage Inequality and Segregation by Skill," NBER working paper W5718.
- Lambson, Val. 1991. "Industry Evolution with Sunk Costs and Uncertain Market Conditions," *Int. J. Indust. Org.* 9:2, pp. 171-96.
- Levinsohn, James. 1993. "Testing the Imports-as-Market-Discipline Hypothesis," *J. Int. Econ.* 35:1-2, pp. 1-22.
- Levinsohn, James and Amil Petrin. 1999. "When Industries Become More Productive, Do Firms? Investigating Productivity Dynamics," NBER working paper W6893.
- Lichtenberg, Frank R. 1992. *Corporate Takeovers and Productivity*. Cambridge: MIT Press.
- . 1995. "The Output Contributions of Computer Equipment and Personnel: A Firm Level Analysis," *Econ. Innovation New Tech.* 3:3-4, pp. 201-17.
- Lichtenberg, Frank R. and Donald Siegel. 1990a. "The Effects of Leveraged Buyouts on Productivity and Related Aspects of Firm Behavior," *J. Finan. Econ.* 27:1, pp. 165-94.
- . 1990b. "The Effects of Ownership Change

- on the Employment and Wages of Central Office and Other Personnel," *J. Law Econ.* 33:2, pp. 383–408.
- . 1991. "The Impact of R&D Investment on Productivity—New Evidence Using Linked R&D–LRD Data," *Econ. Inquiry*, 29:2, pp. 203–29.
- Lucas, Robert E. 1978. "On the Size Distribution of Business Firms," *Bell J. Econ.* 9:2, pp. 508–23.
- McGuckin, Robert H. and Sang V. Nguyen. 1995. "On Productivity and Plant Ownership Change: New Evidence from the LRD," *RAND J. Econ.* 26:2, pp. 257–76.
- McGuckin, Robert H. and George Pascoe. 1988. "The Longitudinal Research Database: Status and Research Possibilities," *Surv. Curr. Bus.* 68:11, pp. 30–37.
- McGuckin, Robert H.; Mary Streitwieser, and Mark Doms. 1998. "Advanced Technology Usage and Productivity Growth," *Econ. Innovation New Tech.* 7:1, pp. 1–26.
- McKinsey Global Institute. 1993. *Manufacturing Productivity*. Washington, DC: McKinsey Co. Inc.
- Moulton, Brent. 1990. "An Illustration of a Pitfall in Estimating the Effects of Aggregate Variables on Micro Unit," *Rev. Econ. Statist.* 72:2, pp. 334–38.
- Nadiri, M. Ishaq. 1970. "Some Approaches to the Theory and Measurement of Total Factor Productivity: A Survey," *J. Econ. Lit.* 8:4, pp. 1137–77.
- Nelson, Richard R. 1981. "Research on Productivity Growth and Productivity Differences: Dead Ends and New Departures," *J. Econ. Lit.* 19:3, pp. 1029–64.
- Olley, G. Steven and Ariel Pakes. 1996. "The Dynamics of Productivity in the Telecommunications Equipment Industry," *Econometrica*, 64:6, pp. 1263–97.
- Oulton, Nicholas. 1998. "Competition and the Dispersion of Labour Productivity amongst UK Companies," *Oxford Econ. Pap.* 50:1, pp. 23–38.
- Pakes, Ariel and Paul McGuire. 1994. "Computing Markov-Perfect Nash Equilibria: Numerical Implications of a Dynamic Differentiated Product Model," *RAND J. Econ.* 25:4, pp. 555–89.
- Pakes, Ariel and Richard Ericson. 1998. "Empirical Implications of Alternative Models of Firm Dynamics," *J. Econ. Theory*, 79:1, pp. 1–45.
- Pavcnik, Nina. 1998. "Trade Liberalization, Exit, and Productivity Improvements: Evidence from Chilean Plants," Princeton U. mimeo.
- Power, Laura. 1998. "The Missing Link: Technology, Investment, and Productivity," *Rev. Econ. Statist.* 80:2, pp. 300–13.
- Salter, W.E.G. 1960. *Productivity and Technical Change*, Cambridge U. Press, Cambridge.
- Tybout, James R. 1993. "Internal Returns to Scale as a Source of Comparative Advantage: The Evidence," *Amer. Econ. Rev.* 83:5, pp. 440–44.
- . 2000. "Manufacturing Firms in Developing Countries: How Well Do They Do, and Why?" *J. Econ. Lit.* 38:1, pp. 7–40.
- Tybout, James R. and Daniel M. Westbrook. 1996. "Trade Liberalization and the Dimensions of Efficiency Change in Mexican Manufacturing Industries" in *Sources of Productivity Growth*. David Mayes, ed. Cambridge: Cambridge U. Press.
- Womack, James P.; Daniel T. Jones, and Daniel Roos. 1989. *The Machine That Changed the World*. Cambridge: MIT Press.