



Information Systems Research

Publication details, including instructions for authors and subscription information:
<http://pubsonline.informs.org>

The Productivity of Information Technology Investments: New Evidence from IT Labor Data

Prasanna Tambe, Lorin M. Hitt,

To cite this article:

Prasanna Tambe, Lorin M. Hitt, (2012) The Productivity of Information Technology Investments: New Evidence from IT Labor Data. Information Systems Research 23(3-part-1):599-617. <http://dx.doi.org/10.1287/isre.1110.0398>

Full terms and conditions of use: <http://pubsonline.informs.org/page/terms-and-conditions>

This article may be used only for the purposes of research, teaching, and/or private study. Commercial use or systematic downloading (by robots or other automatic processes) is prohibited without explicit Publisher approval, unless otherwise noted. For more information, contact permissions@informs.org.

The Publisher does not warrant or guarantee the article's accuracy, completeness, merchantability, fitness for a particular purpose, or non-infringement. Descriptions of, or references to, products or publications, or inclusion of an advertisement in this article, neither constitutes nor implies a guarantee, endorsement, or support of claims made of that product, publication, or service.

Copyright © 2012, INFORMS

Please scroll down for article—it is on subsequent pages



INFORMS is the largest professional society in the world for professionals in the fields of operations research, management science, and analytics.

For more information on INFORMS, its publications, membership, or meetings visit <http://www.informs.org>

The Productivity of Information Technology Investments: New Evidence from IT Labor Data

Prasanna Tambe

Stern School of Business, New York University, New York, New York 10012,
ptambe@stern.nyu.edu

Lorin M. Hitt

The Wharton School, University of Pennsylvania, Philadelphia, Pennsylvania 19104,
lhitt@wharton.upenn.edu

This paper uses newly collected panel data that allow for significant improvements in the measurement and modeling of information technology (IT) productivity to address some longstanding empirical limitations in the IT business value literature. First, we show that using generalized method of moments–based estimators to account for the endogeneity of IT spending produces coefficient estimates that are only about 10% lower than unadjusted estimates, suggesting that the effects of endogeneity on IT productivity estimates may be relatively small. Second, analysis of the expanded panel suggests that (a) IT returns are substantially lower in midsize firms than in Fortune 500 firms; (b) they materialize more slowly in large firms—in midsize firms, unlike in larger firms, the short-run contribution of IT to output is similar to the long-run output contribution; and (c) the measured marginal product of IT spending is higher from 2000 to 2006 than in any previous period, suggesting that firms, and especially large firms, have been continuing to develop new, valuable IT-enabled business process innovations. Furthermore, we show that the productivity of IT investments is higher in manufacturing sectors and that our productivity results are robust to controls for IT labor quality and outsourcing levels.

Key words: business value of IT; economics of IS; econometrics; productivity; IT labor

History: Anitesh Barua, Senior Editor; Ravi Bapna, Associate Editor. This paper was received on July 1, 2010, and was with the authors 4 months for 2 revisions. Published online in *Articles in Advance* March 7, 2012.

1. Introduction

This paper uses newly collected data to analyze the productivity of information technology (IT) investments in a large sample of firms through 2006. Over the last 15 years, there has been considerable progress in the literature linking information technology investment to organizational performance, driven by the availability of large-sample, firm-level data on information technology capital and, to a lesser extent, on complementary organizational practices (Lichtenberg 1995, Brynjolfsson and Hitt 1996, Dewan and Min 1997, Bresnahan et al. 2002). These papers show that IT generates significant returns, typically exceeding the cost of IT capital, and that firms with certain complementary organizational practices realize greater returns from their IT investments (see Melville et al. 2004, Stiroh 2004, or Brynjolfsson and Saunders 2010b for comprehensive surveys). Despite this considerable progress, these analyses have some significant limitations that can be addressed through innovation in data collection and the improved methods that additional data can enable.

In this paper, we use new panel data to address three specific limitations of prior work on IT produc-

tivity. First, a persistent concern in the IT value literature has been establishing how much of the excess rate of return observed for IT investment is because of reverse causality or the endogeneity of IT investment (Lee et al. 1997, Brynjolfsson and Hitt 2000, Aral et al. 2006). Although many prior studies have addressed this concern by using instrumental variables (IV) techniques, the lack of good instruments for predicting firm-specific IT investments leads to high estimation variance. Modern methods have been developed for addressing this problem,¹ but they do not perform well on existing IT data sets.

The second and third limitations relate to sample composition. Most prior work has been restricted to the analysis of large firms (e.g., Fortune 1000) from the mid-1980s to the mid-1990s because that is the sample frame in available data. Although this work has done much to dispel the so-called “productivity paradox,” we know relatively little about whether the

¹ These methods utilize the time dimension of the panel to provide instruments (Arellano and Bond 1991) or use the time-series behavior of other inputs to make corrections for estimates of capital coefficients (Levinsohn and Petrin 2003, Olley and Pakes 1996, Akerberg et al. 2006).

pattern of IT returns observed in large firms generalizes to smaller U.S. firms (see Dedrick et al. 2003 for a similar claim) because firms of different sizes might differ in their ability to assimilate new IT investments or may have made different levels of investment in complementary IT or organizational practices in the past. Furthermore, the time period most extensively studied in the past (~1987 to ~1997) was characterized by extensive organizational transformation and predates the large boom in computer investment that occurred in the United States in the late 1990s, leaving open the question of whether the productivity of IT investments has changed materially since the “Internet revolution.” Higher returns to IT investment in recent years would suggest that firms are continuing to develop IT-enabled process innovations, but economists are increasingly concerned that this recent period might be characterized by declining returns to IT investment, suggesting that the stock of potential IT-enabled business innovations is being depleted (e.g., Stiroh 2008).

These gaps in our understanding of IT productivity have persisted because existing firm-level IT research has generally had to rely on one of four possible data sources. The *Computerworld* (Brynjolfsson and Hitt 1996) and *InformationWeek* data sets (Lichtenberg 1995) rely on annual surveys of large firms. Although for a number of years these were the only IT data available at the firm level, they have relatively small cross sections (200–300 firms/year), and despite a consistent sample frame, there is limited year-to-year consistency in firm responses, making them unsuitable for panel data methods. Prior to 2003, the U.S. Census Bureau collected data on IT expenditures through various special surveys (e.g., the 1999 Computer Network Use Supplement), but these surveys, although broad and highly detailed, are not consistently available over time. Since 2003, the Census has expanded the Annual Capital Expenditure Survey (ACES) to include questions on hardware and software IT expenditures, which currently yields a five-year panel. This is likely to be a valuable resource in the future, especially for firm size or industry comparisons, but the current panel has a limited time dimension and consequently cannot be compared to the best available data from other sources that cover prior periods.

The most comprehensive data set is the Computer Intelligence Technology Database (CITDB), which is a panel of large firms (principally Fortune 1000) from roughly 1987 to the present. However, in 1994, the CITDB changed its method of capital valuation and by 1996 no longer provided estimates of computer capital stock at all. Researchers have extrapolated these data to enable IT stock to be calculated approximately through the year 2000, although this is likely

to introduce considerable error, and we are aware of no attempts to extrapolate these data beyond 2000 for productivity calculations.² Consequently, there is currently no existing data set that has good time-series comparability (like the CITDB data), has a long history, is available in current periods, and covers a broad cross section of large and smaller firms.

In this study, we develop a new data set based on IT personnel counts and matching production inputs for approximately 1,800 firms across 20 years (36,000 firm-years from 1987 to 2006). This makes this data source, to the best of our knowledge, one of the more complete firm-level IT panels that has been available to researchers. As we demonstrate below, our data provide a much larger and more recent sample for IT productivity work while retaining the most useful characteristics of existing data such as a consistency in the within-firm time dimension. Although our study is not the first to propose the use of IT labor as a measure of IT investment (see, e.g., Lichtenberg 1995 and Brynjolfsson and Hitt 1996), our data are unique in their scope and consistency, which enables the application of estimators that require a longer time dimension than have been available in other sources of IT panel data.

We first use comparisons between our data and the best available prior data set (CITDB) to show that we are able to replicate prior results, and we use the CITDB data as an instrument for our IT labor measures to demonstrate that measurement error in our data is sufficiently low. Describing the measurement error properties of our data is important for establishing the utility and accuracy of these data and to compare our measures to the best available alternatives. Next, we use these new IT measures to address the three limitations of prior work identified earlier. Our first major finding is that endogeneity, at least in the form addressed in modern microproductivity measures, does not substantively affect current IT estimates, nor, likely, prior IT estimates. Estimates that address endogeneity only lower measured IT elasticity by 10% versus the methods used in prior work. Furthermore, measuring IT using labor, which is likely to be especially subject to endogeneity bias, suggests that the bias in capital-related studies is even smaller. We also show that high IT returns are not attributable to some other potentially important sources of endogeneity, unobserved differences in labor quality, or outsourcing.

²Chwelos et al. (2010) provide a method for extending CITDB 1994 valuation data through 1998 by imputing the values of equipment in the earlier part of the data set and adjusting for aggregate price changes. However, this differs from the method employed by CITDB, which determined equipment market values by looking at actual prices in the new, rental, and resale computer markets, and cannot be reasonably applied to more recent data because of the substantial time lag from the data used to calibrate the models.

Second, we find that both large and midsize firms make similar investments in IT relative to their size. However, larger firms appear to realize greater marginal product from these investments, whereas midsize firms experience the benefits of these investments much more quickly. These observations are consistent with the argument that adjustment costs are lower for IT investments in smaller firms, but larger firms are better positioned to take advantage of IT-related complements. Third, the productivity effects of IT investment have persisted over time—returns to IT spending continued to be higher for large firms through the late 1990s and into the current decade. In fact, we provide evidence that the measured marginal product of IT labor is higher in recent years than in the past in firms of all sizes, in contrast to recent work that suggests that the link between IT spending and productivity may have changed materially since 2000 (Stiroh 2008, Jorgenson et al. 2008).

2. Background

2.1. The Productivity Estimation Framework

The contribution of information technology to productivity has most commonly been determined using methods from production economics, which allow researchers to estimate the relationship between various production inputs, such as capital and labor, and firm output. This literature relies on the concept of a *production function*, an econometric model of how firms convert inputs to outputs. Economic theory places some constraints on the functional form used to relate these inputs to outputs, but several different functional forms are widely used depending on the firm's economic circumstances. Perhaps the most widely used of these forms is the Cobb–Douglas specification. Aside from being among the simplest functional forms, this specification has the added advantage that it has by far been the most commonly used model in research relating inputs such as information technology to output growth at a variety of levels of aggregation (e.g., plant, firm, industry) and forms the basis for productivity measurement of the U.S. economy as a whole. Moreover, because Cobb–Douglas can be considered a first-order approximation of an arbitrary production function, it is well suited to estimating the contribution of inputs to outputs, which are typically quoted at the sample mean, the region where a first-order approximation is especially accurate.³ In this study, therefore, we assume

that firms produce output via a Cobb–Douglas production function whose inputs are capital (K), non-IT labor (L_0), and IT labor (L_1). We have chosen to use value-added (VA) as the dependent variable for consistency with most prior IT-value research, although our results are similar when we utilize gross output as the dependent variable and incorporate materials as an additional covariate. Although most production function estimates do not distinguish between IT labor and non-IT labor, making this distinction allows us to separate the contribution of IT workers. An estimable model of the Cobb–Douglas function, similar to the models used in both Lichtenberg (1995) and Brynjolfsson and Hitt (1996), can be written as

$$\ln VA = \alpha_K \ln K + \alpha_L \ln L_0 + \alpha_{IT} \ln L_1, \quad (1)$$

where L_1 is IT labor, L_0 is all other labor, and the model can be estimated using standard regression techniques such as ordinary least squares (with suitable standard error corrections for panel data) or panel methods such as fixed effects. The coefficient estimate on the IT labor input (α_{IT}) is the *output elasticity* of information technology labor, the percentage increase in output generated by a 1% increase in the IT labor input. The output elasticity has the advantage that it is independent of the units used to measure the inputs and outputs, but it has the disadvantage that it cannot be easily compared across different samples that have different average levels of IT investment or other factor input shares. Therefore, in addition to reporting output elasticities, we compare the results from estimates across different regressions by computing the *marginal product* (MP), the amount of additional output that can be produced for an additional unit, such as a dollar, of a given input.⁴ In equilibrium (under textbook assumptions such as in Varian 1992, Chapter 2), a profit-maximizing firm should invest in an input until the marginal product is equal to the marginal cost. Moreover, because firms are likely to invest in the highest value uses of any input first, average returns to all input units are likely to equal or exceed the marginal product. A marginal product in excess of marginal cost indicates an “excess return,” which may be reconciled with profit maximization in a variety of ways, including factor adjustment costs or unmeasured complementary inputs (see extensive discussions of this issue in Brynjolfsson and Hitt 2003 or Brynjolfsson et al. 2002). Regardless of the reason, inputs with measured excess return not only contribute to increased output but also to growth in multifactor productivity.

³ Estimates of transcendental logarithmic or constant elasticity of substitution production functions using these data (methods used in prior work such as Dewan and Min 1997 or Brynjolfsson and Hitt 1996) yield nearly identical estimates of the output elasticities of all inputs as our Cobb–Douglas estimates, as expected.

⁴ The marginal product for computers is computed as the output elasticity multiplied by the ratio of output to computer input (the reciprocal of the factor input share).

2.2. Limitations and Related Literature

The production function approach described above has formed the basis of a large and influential empirical literature on IT productivity. Although recent IT value research has used this framework to answer questions that go beyond estimating private returns to IT investment (e.g., Cheng and Nault 2007, Tambe and Hitt 2012), several open questions remain about the estimates that have been produced using this approach. The primary concern of the production function approach is that unobserved demand or productivity shocks may induce firms to make greater IT investments, creating a reverse causal relationship between IT investment and output and creating an upward bias in elasticity estimates. The most common approach to dealing with this issue is the use of instrumental variables, but good instruments for IT investment have been difficult to find. Recent work has used organizational survey responses from limited numbers of firms to create instruments based on adjustment cost differences (Tambe et al. 2011), but finding effective instruments for IT investment that can be applied to large samples of firms remains a persistent issue in the IT productivity literature (Brynjolfsson and Hitt 2000).

There has been recent progress, however, in the development of new techniques for the econometric identification of production functions related to the use of (1) dynamic panel data estimators (Blundell and Bond 2000) and (2) estimators that use structural modeling techniques to identify production functions (Olley and Pakes 1996, Levinsohn and Petrin 2003, Akerberg et al. 2006). A challenge with these classes of estimators, however, is they require more from the data—for instance, dynamic panel data estimators generally require longer panels than have historically been available for firm-level IT data, and structural estimators also place demands on the time dimension because they utilize change in noncapital inputs to identify endogeneity biases. The application of these estimators, however, can resolve many of the classical endogeneity concerns related to the effects of unobserved inputs on production function estimates.

The second major shortcoming of prior work in this stream is related to sample restrictions—most existing work has been restricted to large firms before the mid-1990s. A central finding from the IT productivity literature is that IT returns in Fortune 1000 firms were growing larger over longer time differences from the mid-1980s to mid-1990s, suggesting that firms were taking time to reorganize work processes and build the necessary organizational complements (Brynjolfsson and Hitt 2003, Brynjolfsson et al. 2002). Some examples of adjustment costs borne by firms when installing new IT are the redesign of jobs and work routines, the creation of new incentive systems, the development

of new software applications, and retraining employees in how to use IT-enabled systems (Applegate et al. 1988, Bresnahan and Greenstein 1997).

However, the experiences of the larger firms upon which these findings are based may not reflect the experiences of smaller firms (Dedrick et al. 2003). First, prior work ties adjustment cost structure to the specificity of complementary assets and argues that large firms tend to face greater adjustment costs (such as those associated with custom application software) because of idiosyncratic work processes and the tacit organizational knowledge that accumulates in larger firms (Ito 1995). Smaller firms can more easily use common, standardized applications because they have less firm specificity embedded in their internal firm transactions. Large firms, therefore, should face greater disruptions and longer productivity lags when reorganizing but may also eventually build more valuable organizational assets than would smaller firms because the firm specificity embedded in the organizational assets of large firms may make them more difficult for other firms to imitate.

Furthermore, the adoption of information technologies requiring extensive business process reorganization may be disproportionately costly for large firms (McElheran 2010). This argument is consistent with a strategy and economics literature that demonstrates that radical innovations reduce the value of a firm's existing know-how, making large firms, which are more likely to have made greater investments in complex routines, less likely to adopt disruptive innovations (Arrow 1962, Tushman and Anderson 1986, Henderson 1993). However, the reorganization of complex business processes around new information technologies may produce greater long-run benefits for large firms because changes embedded in complex routines are more difficult for competitors to imitate (Mata et al. 1995, Melville et al. 2004). Differences in adjustment costs and appropriability between large and smaller firms, therefore, may correspond to differences in long-run returns to IT investment. Conducting this empirical comparison requires data on a smaller class of firms than has been available through traditional IT data sources.

In addition to understanding how IT benefits differ across firms of different sizes, it is also important to understand how they evolve across time to obtain a better understanding of the relationship between IT and productivity growth after the mid-1990s. Whether IT continued to contribute to productivity in the last decade depends in part on whether firms continued to grow IT-related intangible assets or whether the restructuring observed through the mid-1990s marked the end of the adjustment period for firms. Nevertheless, there has to date been little work analyzing how differences in adjustment costs across

firms or time might affect productivity or growth rates.⁵ Addressing these questions requires a sample that not only includes midsized firms but also is long enough through which to observe temporal differences in IT returns.

Sample size issues in prior studies also limit the ability to subdivide the data for making sectoral comparisons. Different industries rely on different production technologies and different organizational practices. As a general-purpose technology, IT is likely to provide some benefit to all industries, but little is known about how these benefits may vary, and some of this variation may be of significant economic interest. For instance, it has been argued that the contribution of information technology value and its appropriability may be significantly different in service industries (e.g., Steiner and Teixeira 1990, Roach 1991) and that measurement issues may obscure the value being created (see, e.g., Brynjolfsson 1993). Others have argued that the value of information technology has been disproportionately captured by computer-producing rather than computer-using industries (Gordon 2000), which is critical for understanding the long-term benefits of computer investment because it implies productivity gains are driven from pure technological progress (e.g., Moore’s law) rather than a combination of technological progress and complementary organizational innovation.

3. IT Measure Validation

3.1. IT Labor Measures

To resolve the data limitations described above, we develop new firm-level IT measures that can be incorporated into the production function framework described above to estimate the contributions of IT to productivity. We obtain measures of firm-level IT employee counts from data based on the employment histories of a large sample of information technology workers. These employment histories, of the type shown in Table 1, were collected through a partnership with a leading online job-search website and include information for each worker on employer name, job title, and dates of employment for every position ever held by that worker. Employment histories at this website have been posted by close to 10 million unique individuals who are passively or actively seeking jobs. In addition to posting full résumé data, visitors manually enter firm names, job titles, and dates associated with all previous employment spells as well as some human capital data such

⁵ A notable exception is Bloom et al. (2008), who argue that higher IT-related organizational adjustment costs for firms in Europe (created by differences in labor regulations) delayed IT investment and lowered returns to IT investment in European firms relative to firms in the United States.

Table 1 Employee History Database

Employee	Employer name	Job title	Start date	End date
IT worker	Firm name 3	Project manager	5-01-2006	Present
IT worker	Firm name 2	Software engineer	9-01-2003	3-15-2006
IS manager	Firm name 2	Director of technology	4-01-2006	Present
IS manager	Firm name 1	MIS manager	1-01-2006	3-20-2006

as level of education and managerial experience. They also select occupational categories, such as information technology, sales, finance, or production.

In this study, we specifically use information on the employment histories of the information technology workers to construct our measures of IT labor as well as on the employment of all workers to account for sampling issues. About 600,000 of the workers in our data set reported information technology as their primary occupation in 2006, representing perhaps as much as 10%–15% of the total U.S.-based IT workforce.⁶ Table 2 shows some summary statistics comparing the IT workers in our sample to the sample of workers in the Current Population Survey (CPS) who list information technology as their primary occupation. Although the workers in our sample tend to have shorter job tenures, the educational distribution is similar across the two samples. Workers report their employer name for each employment spell, which we standardize using string-matching techniques and then match against a number of firm name databases, such as the Compustat database, the National Bureau of Economic Research (NBER) Patents database, and a privately developed database that contains a large number of common name permutations of publicly traded firms.⁷ These methods are similar to processes used for other productivity-oriented data sets (the NBER Patents Database or the Census Longitudinal Employer-Household Dynamics (LEHD)). Although incorrect or missing matches because of unusual spelling or name structure may introduce error into our IT measures, the errors are likely to be random and can be addressed by comparing our measures to those produced by the CITDB. In addition, because we correct for sampling biases (see below) many of these errors will cancel out unless IT workers consistently spell their firm names differently than other workers do at the same firm.

⁶ Based on the Bureau of Labor Statistics (BLS) report of the total number of workers employed in “Computer and Mathematical Occupations” in 2006.

⁷ This alternative data source was obtained from a separate online service in which individuals post their employment histories. Unlike our primary source, participants post both the ticker symbol and an employer name. These data provide a large number of permutations of company names that can be linked to a common identifier.

Table 2 Demographic Statistics for Information Technology Worker Sample


Education	Matched résumé sample	CPS 2006
High school degree or less	24.7	25.1
Vocational degree	2.8	0.8
Two-year degree	14.3	10.8
Four-year degree	38.8	42.8
Graduate degree	18.6	18.9
Doctorate	0.7	1.7
Average job tenure	4.00	6.33

Source. Current Population Survey (CPS) 2006.

A single cross section of résumés contains a full, time-series employment history for each individual in the database. To convert our sample of employees to measures of total IT employment for the firms we analyze, we make a number of adjustments to correct for overall, firm-specific, and time-varying sampling rates. Our primary assumption is that the workers we observe are “sampled” from the underlying population of all workers. The number of IT workers reporting employment in a particular firm represents a sample proportion $\hat{p}_j = x_j/N$, where N is the total number of samples and x_j is the number of workers reporting employment at firm j . If these employment histories were randomly drawn from the population of IT workers, the sample proportion would be an unbiased estimate of the true proportion of IT workers employed by firm j . However, employees cannot be viewed as having been randomly sampled from the population. Instead, the likelihood that a worker posts an employment history is influenced by firm-level factors, such as the bankruptcy or financial decline of an employer, and occupation-specific factors, such as average education levels within an occupation, creating differences in the underlying sampling rate among both occupations and firms.

To account for these underlying sampling differences and recover firm-level IT employment levels from our raw IT employment sample estimates, we assume that firm- and occupation-specific factors that affect the likelihood of posting one's employment history are uncorrelated.⁸ The firm-specific sampling

⁸This assumption is violated if idiosyncratic factors at particular firms affect IT worker turnover. Then the IT employment measure will be too high when events force an abnormally high number of IT workers to exit the firm and too low when IT workers are retained at an abnormal rate. This biases our IT productivity estimates upwards if unobserved factors are associated with both a high IT turnover rate and higher productivity levels. However, (1) we control for IT outsourcing in our analysis, which is one of the primary reasons that we might see this pattern, and (2) if this were a significant source of error, it would be captured by our measurement error analysis and our correlation tests below because IT capital would be uncorrelated with this source of error.

Figure 1 Estimated Average Employee Sampling Rate by Year

rate (θ_j) can then be estimated by comparing the number of workers in all occupations in our data from a particular firm (x_j) against total employment at that firm (L_j) as reported in Compustat:

$$\hat{\theta}_j = \frac{x_j}{L_j}. \quad (2)$$

Estimates of the number of IT workers employed by the firm can be adjusted by scaling the number of sampled IT workers employed by the firm by the firm-specific sampling rate, which accounts for any firm-specific factors that may affect an employee's propensity to appear in our sample. This division process also reduces the error due to name mismatches.

Because of the entry and exit of workers from the workforce over time, we will naturally capture a larger fraction of the workforce in more recent time periods. The average sampling rate, shown in Figure 1, declines from 12% in 2006 to about 4% in 1987. The dip in our sampling rate from 2005 to 2006 is because we count only employed workers and unemployment rates in our sample rose over this period. Regardless, even the 4% sample appears sufficient for relatively precise estimates of the overall IT workforce, and because of the large number of samples, the error variance due to sampling in this measure is likely to be small.⁹ However, later in this analysis we use instrumental variables to test how errors-in-variables impacts the estimates produced by these data and show that the variance from all sources of random error (sampling, matching, etc.) is relatively small.

The final IT employment data set includes measures for more than 36,000 firm-years from 1987 to 2006. We match these data to Compustat to compute other production inputs. We compute *value-added* as sales minus nonlabor variable costs deflated by an output deflator at the two-digit Standard Industrial

⁹If workers in our sample are randomly drawn from the population, the sampling error variance is $p_j(1-p_j)/N$, where p_j is the true proportion of IT workers employed by firm j and N is the overall number of samples.

Table 3 Means and Standard Deviations for Full Sample

Variable	Mean	Std. dev.	<i>N</i>
<i>Value added</i> (\$ million)	1,073.7	3,402.6	36,305
<i>Output</i> (\$ million)	3,135.4	10,957.5	36,305
<i>Capital</i> (\$ million)	2,106.4	8,191.6	36,305
<i>Materials</i> (\$ million)	2,061.7	8,187.8	36,305
<i>Non-IT employees</i> (×1,000)	13.8	43.5	36,305
<i>IT employees</i>	277.4	1,035.6	36,305
<i>IT capital</i> (CITDB) (\$ million)	24.0	55.6	4,733
<i>% IT outsourcing</i>	14.8	17.6	237

Classification (SIC) level (see, e.g., Brynjolfsson and Hitt 2003, Bureau of Economic Analysis¹⁰). *Capital* is computed from the gross book value of capital converted to current dollars based on average capital age and an industry-specific deflator, which in turn is calculated by dividing total depreciation by current depreciation and averaging this figure over three years—see Hall (1990) and the deflators from the Bureau of Economic Analysis fixed asset tables.¹¹ Number of *employees* and *industry* is taken directly from Compustat. Non-IT employees are computed as total employees minus IT employees. Overall, our sample covers 57% of all publicly traded firms with \$10 million or more in sales that report employees and capital stock, making this the broadest and most comprehensive sample of firms used in IT productivity research.¹² The typical firm in our sample is large, with an average value-added of about \$1.15 billion (see Table 3). The average number of IT workers employed at a firm in our sample is about 277.

3.2. IT Labor Measure Benchmarking

Beyond theoretical sampling considerations, we can test whether these IT measures are a good representation of firm-level IT inputs by comparing them against external data sources. First, we collect IT data sets from different years, some of which have been used in earlier IT productivity studies. These data sources (primarily survey based) focus on IT inputs in larger firms. High correlations between our measures and the measures used in prior surveys suggest that we are capturing the same underlying constructs and regressions using these measures should yield comparable results to those reported in earlier studies.

¹⁰ See http://www.bea.gov/industry/gdpbyind_data.htm (Accessed February 2, 2012).

¹¹ See <http://www.bea.gov/national/FA2004/index.asp> (Accessed February 2, 2012).

¹² Employees, sales, and capital stock are the minimum variables for estimating production functions. Screening firms by sales also helps eliminate nonoperating holding companies. The \$10 million sales cutoff was chosen to represent firms likely to have at least five IT employees (at a fully loaded cost of \$100,000 per IT employee and a 5% IT/sales ratio), which is close to the smallest number detectable in our data set given our approximately one in seven IS employees sampling rate.

Statistics from the comparisons are shown in Table 4. Column (1) compares our measures against 1988–1992 *Computerworld* survey data in which respondents reported levels of aggregate IT labor expense (these data were used for Brynjolfsson and Hitt 1996). The correlation between our employment measures and the *Computerworld* labor expense measures is 0.63, and the means from the two data sets are consistent with a reasonable level of expenditure per IT employee. Column (2) compares our labor-based measures against the CITDB capital stock measures, available from 1987 to 2000 (e.g., see Brynjolfsson and Hitt 2003). The correlation between these two sets of measures is 0.57. In Columns (3) and (4) we benchmark our measures against surveys that explicitly ask for the number of IT employees within the organization, allowing for a more direct comparison of our measures. Column (3) compares our measures against 1995–1996 *Informationweek* data, and Column (4) compares our measures against a 2001 survey conducted by researchers at the Massachusetts Institute of Technology (MIT) (e.g., see Tambe et al. 2011). Our measures appear to be highly correlated with both sets of survey data. Finally, in Column (5) we compare our IT labor data from 2006 with IT employment numbers collected from a large sample of firms in 2008. The simple correlation between the two sets of measures is 0.78. Overall, these comparisons suggest that (1) our measures are closely correlated with external sources of firm-level IT data and therefore representative of IT expenditures at the firm level and (2) these correlations extend throughout the duration of our panel.

Figure 2 shows how our estimates compare to overall national IT employment from 1998 to 2006 where our parameter estimates are set so the values are equal in 2000 but can differ in all other years.¹³ The trend lines suggest that the entry and exit of IT workers in our data by year is in proportion to overall levels, and the two series show a correlation of 0.77.

Table 5 shows the distribution of these observations by one-digit SIC industry. An apparent feature of these data is that they include a large number of service sector observations, which has been a notable limitation of much of the e-commerce data collected by U.S. statistical agencies (Atrostic et al. 2001). In Table 6, we examine statistics regarding the sizes of firms in our sample. Fortune 500 firms form only about 29% of our sample. On the other hand, the firms in the matched CITDB sample have employment and total sales statistics that are very close to those of Fortune 500 firms. The remaining firms in our sample have a headcount that is three to four

¹³ The BLS only provides national employment numbers by occupation from 1998 onward.


Table 4 Comparison of Résumé-Based Measures with External IT Measures

	(1)	(2)	(3)	(4)	(5)
External IT sample	<i>Computerworld</i> ^a	CITDB ^b	<i>InformationWeek</i> ^c	MIT ^c	MIT ^d
Years	1988–1992	1987–2000	1994–1995	2001	2006 ^e
Coefficient of Variation					
External sample	0.42	2.32	3.33	2.76	1.95
Matched sample	1.16	2.50	1.91	2.06	1.78
Correlations					
Correlation	0.58	0.56	0.52	0.63	0.78
Correlation of logs	0.53	0.62	0.55	0.74	0.75
Spearman	0.63	0.58	0.12	0.72	0.75
Firm size controls	0.37	0.44	0.49	0.53	0.59
<i>N</i>	600	4,754	283	80	164

Note. All correlations with multiyear samples include year dummies.

^aMeasured in millions of dollars of IT labor expenses (*Computerworld* surveys). ^bMeasured in millions of dollars of IT capital stock (CITDB data). ^cMeasured in number of IT employees (*InformationWeek* and MIT surveys).

^dMeasured in number of IT employees in 2008 (MIT survey). ^eCompared with 2008 values of MIT survey sample.

Figure 2 Estimated National Employment of Computer Workers Compared with Bureau of Labor Statistics (BLS) Computer Worker Employment Data

times smaller, so there is a substantial size difference between the firms in our sample and the firms that have been studied in prior IT productivity research.

The size distribution of the firms in our data is presented in Table 7 and is compared against (a) the business size distribution data of all U.S. firms collected by the U.S. Census Bureau and (b) the size distribution of firms in the CITDB data. The Census data on size distribution are collected from the Statistics of U.S. Businesses Data collected by the U.S. Census Bureau.¹⁴ We compare these size distributions in 1998, the most recent year where all three key data sets overlap (ours, CITDB, and Census Bureau), and in 2005, where our data have the

Table 5 Firm-Year Observations by Industry

Major industry	<i>N</i>	Percentage of sample
Construction	37	0.1
Durable manufacturing	110,142	30.7
Nondurable manufacturing	6,383	17.6
Transportation and utilities	3,458	9.5
Wholesale trade	1,506	4.2
Retail trade	3,670	10.1
Financial services	2,345	6.5
Nonfinancial services	7,764	21.4
Total	36,305	100

maximum sampling rate and overlaps with Census. A comparison of the distributions indicates that both our IT labor data and the CITDB data oversample large firms relative to firms of other sizes. However, this is at least partly because we restrict our sample to public firms for which other production inputs are available. An important difference between the two data sets is that the IT labor data offer significant coverage for firms with fewer than 10,000 employees, whereas the CITDB data have few firms in these categories. Therefore, although limited to public firms and not representative of the full size distribution of the U.S. economy, this data set offers a large enough number of observations across the distribution to make meaningful statistical inferences about the effects of IT investment in firm size categories outside the Fortune 1000.

3.3. Comparison with CITDB Data

In this section, we perform a more rigorous comparison of our IT labor data set with the CITDB capi-

¹⁴ Information available at <http://www.census.gov/econ/susb/> (Accessed February 2, 2012).

Table 6 Output and Employment Comparisons Between Fortune 500 Firms and All Other Firms

	Full IT labor sample	Fortune 500	CITDB matched	500 Non-Fortune
Value added (\$ million)	1,074.0	2,820.3	2,746.6	541.5
Output	3,135.4	7,168.7	6,594.1	1,487.0
IT employment (\$ million)	288.9	657.8	635.7	176.5
Employment	13,831.2	32,298.9	33,785.6	6,283.5
N	36,305	10,533	4,754	25,772

Table 7 Firm Size Distributions

Data source	1998 SUSB	1998 CITDB	1998 IT labor	2005 SUSB	2005 IT labor
<500 employees	99.71%	0.45%	23.74%	99.71%	20.51%
500–999 employees	0.14%	0.23%	10.18%	0.15%	10.69%
1,000–2,499 employees	0.09%	3.40%	17.60%	0.09%	16.78%
2,500–4,999 employees	0.03%	8.62%	13.73%	0.03%	14.22%
5,000–9,999 employees	0.02%	20.63%	12.29%	0.02%	13.16%
10,000 or more employees	0.02%	66.67%	22.47%	0.02%	24.64%
Total	5,579,177	441	2,279	5,983,546	2,545

Notes. Cell numbers represent percentage of total firms in each size class category. SUSB is the Statistics of U.S. Businesses Data collected by the U.S. Census Bureau. Information available at <http://www.census.gov/econ/susb/>. Numbers for SUSB firms are firms with a nonzero number of employees.

tal stock data set, the largest of the firm-level data sets that have appeared in the IT productivity literature (Brynjolfsson and Hitt 2003). Our measures differ from the CITDB data in a number of ways that create opportunities for new analysis. First, as described above, they include a larger cross section of firms than do the CITDB capital stock data and have a longer time dimension. Compared to CI, we have approximately 3 times as many firms and 2.5 times as long a time series, so we can compute meaningful statistical tests on industry and time period subsamples. Second, it is likely that our measures have less random error variance than the CITDB because of the way the data are constructed, a conjecture which we can test below.¹⁵ Third, there may be advantages of focusing on labor-based rather than capital-based measures of IT for studying recent time periods. On average capital composes only 10%–15% of a firm’s total IT budget or capital stock (Brynjolfsson et al. 2006) and perhaps even less in recent time periods (Saunders 2010b). By contrast, labor expenditures may account for more than 30% of a firm’s IT budget (*InformationWeek* 2006), and the intangible assets produced by IT staff such as internally developed software may be more than five times larger than the stock of computer hardware (Saunders 2010b). Finally, many of the remaining factors that make up the budget, such as training and

process change, may be more closely correlated with levels of IT labor than is capital stock.

To benchmark the performance of our measures, we compare their behavior in production function regressions alongside that of the CITDB capital stock measures. Between 1987 and 2000, our sample includes 4,745 firm-year observations for which we have both CITDB IT capital and our IT employment measures. The simple correlation between the IT measures for the overlapping sample is 0.57, suggesting the measures are closely related but have some independent variance.¹⁶ Column (1) in Table 8 shows cross-sectional estimates using only the CITDB data. All analyses in this table utilize ordinary least squares (OLS) with Huber-White robust standard errors to account for repeated observations of the same firm over time and any conditional heteroskedasticity in the model. In general, Huber-White standard error estimates will be used throughout the paper whenever OLS estimators are used unless otherwise noted. Regression estimates using these data indicate an elasticity of CITDB IT capital of about 0.124 ($t = 8.9$). Elasticity estimates from the labor-based measures using the same sample of firms, shown in Column (2), are somewhat higher at 0.155 ($t = 7.4$). Column (3) includes both capital- and labor-based measures in the same model. The coefficient estimate on capital is 0.102 ($t = 6.8$), and the estimate on labor is 0.121 ($t = 6.1$). These estimates represent reductions of about 20% relative to the coefficients when estimated

¹⁵ See Greenan and Mairesse (2000) for an example of IT productivity measurement when IT stocks are measured using labor measures with sampling error. They demonstrate that it may be possible to accurately estimate IT expenditure using employment data with samples as small as one to five workers per firm. Our samples are considerably larger.

¹⁶ The correlation between the two measures is 0.32 after including size controls and 0.23 after including size and industry controls.

Table 8 Comparisons of IT Employment and IT Capital Stock Measures

DV: Value added	(1)	(2)	(3)
	1987–2000	1987–2000	1987–2000
<i>Non-IT employees</i>		0.491 (0.029)**	0.474 (0.027)**
<i>Capital</i>		0.297 (0.018)**	
<i>Employees</i>	0.571 (0.023)**		
<i>Non-IT capital</i>	0.272 (0.018)**		0.264 (0.018)**
<i>IT capital (CITDB)</i>	0.124 (0.014)**		0.102 (0.015)**
<i>IT employees</i>		0.155 (0.021)**	0.121 (0.020)**
Controls	<i>Industry</i>	<i>Industry</i>	<i>Industry</i>
	<i>Year</i>	<i>Year</i>	<i>Year</i>
<i>N</i>	4,745	4,745	4,745
<i>R</i> ²	0.87	0.87	0.88

Notes. Robust standard errors shown in parentheses. All variables are in logs. Industry controls are included at the one-digit SIC level.

** $p < 0.05$.

alone. Estimates of other factor inputs are affected in predictable ways. The labor elasticity drops when IT workers are included explicitly, whereas the capital measures drop slightly when IT capital is explicitly included. This is consistent with both IT capital and IT labor having a marginal product per unit greater than ordinary capital and labor, respectively. However, because IT is still a relatively small portion of capital and labor, the marginal product estimates of non-IT are consistently close to their theoretical values regardless of whether IT is separated or not. The results of this analysis indicate that these IT labor measures perform reasonably when used instead of or alongside the CITDB data in standard productivity regressions.

In addition to making direct comparisons, we can use the CITDB capital stock measures as an instrument for our observed labor measures (and vice versa). Although this does not address reverse causality, this instrument has the potential to eliminate biases because of measurement error (see a similar approach in Brynjolfsson and Hitt 2003), allowing us to gauge the magnitude of bias because of the measurement error and to explicitly estimate error variance in each data set. The key assumption in this analysis is that measurement error in the two data sets should be uncorrelated, and this assumption is likely to be valid because the two data sets are constructed from completely different information using different methods. Specifically, in the presence of measurement error, the coefficient estimate on the mis-measured input (β^*) will be equal to the true estimate

(β) attenuated by an amount equal to the ratio of the signal variance to the total measure variance:¹⁷

$$\beta^* = \beta \left(\frac{\sigma_x^2}{\sigma_x^2 + \sigma_\varepsilon^2} \right). \quad (3)$$

Therefore, the ratio of the error variance to the total measure variance can be computed using the biased (OLS) and corrected (IV) coefficient estimates:

$$1 - \frac{\beta^*}{\beta} = \left(\frac{\sigma_\varepsilon^2}{\sigma_x^2 + \sigma_\varepsilon^2} \right). \quad (4)$$

In Table 9, we present the results of our measurement error analysis. We first use the IT employment data as an instrument for the CITDB data in (1) and (2), and these estimates suggest that the error variance in the CITDB data between 1987 and 1994 is on the order of 40% of the total measure variance, which is within the range suggested by researchers who have used these data in earlier work (Brynjolfsson and Hitt 2003). In Columns (3) and (4), we report estimates when using our primary employment measures over the same time period, with and without the use of the CITDB instrument, in the limited sample of our IT employment data for which the CITDB data are also available. The estimate on the IT employment variable rises after the application of the instrument, which is consistent with a measurement error interpretation, and the difference in magnitudes implies that the variance of the error term is

¹⁷ For a discussion of the classic errors-in-variables model, see Wooldridge (2002).

Table 9 Using Instrumental Variables to Test Robustness to Measurement Error

	(1)	(2)	(3)	(4)	(5)	(6)
	OLS	IV	OLS	IV	OLS	IV
DV: <i>Value added</i>	1987–1994	1987–1994	1987–1994	1987–1994	1995–1997	1995–1997
<i>IT employment</i>			0.034 (0.016)**	0.040 (0.027)	0.101 (0.028)***	0.105 (0.027)***
<i>CITDB IT capital</i>	0.032 (0.017)*	0.054 (0.025)**				
Controls	<i>Industry</i> <i>Year</i>	<i>Industry</i> <i>Year</i>	<i>Industry</i> <i>Year</i>	<i>Industry</i> <i>Year</i>	<i>Industry</i> <i>Year</i>	<i>Industry</i> <i>Year</i>
<i>N</i>	2,176	2,176	1,863	1,863	2,176	2,176
<i>R</i> ²	0.95	0.95	0.95	0.95	0.94	0.94
Implied measurement error ^a		41%		15%		4%

Notes. All variables are in logs. Industry controls are at the one-digit SIC level. Regression also includes *capital*, *non-IT employment*, *industry*, and *year*.

^aVariance of measurement error is reported as a percentage of total measure variance.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$; dependent variable is *value added*.

about 15% of the total measure variance. The average measurement error in our IT employment sample, therefore, appears to be substantially smaller than that in the CITDB data through 1994. It is noteworthy, however, because our sampling rate increases over time, the measurement error in later periods should be lower than in initial periods. We are limited in the extent to which we can test this hypothesis because of the availability of our instrument, which is available and reliably measured only through 1997. However, in Columns (5) and (6), we estimate measurement error in the IT employment data between 1995 and 1997. The change in magnitudes implies a measurement error term in earlier years that is around 5%, which is consistent with our earlier assertion that measurement error should be falling in our sample over time. These estimates suggest (1) that our IT employment measures contain less measurement

error than do the CITDB data, which have been widely used in a number of notable IT productivity studies; (2) that measurement error is decreasing over time in our sample; and (3) that the magnitude of the measurement error, especially in more recent years, is not large enough to raise serious concerns about the estimates produced using these data. In the next section, we report estimates using the full sample for which our employment measures are available.

4. Productivity Estimates

Our analysis begins by replicating cross-section and panel production function estimates using our new data that can be compared to prior work. We then apply new estimators to address the impact of endogeneity. Finally, we consider subsamples to address how IT returns differ across time and firm size.

Table 10 Baseline Estimates Using IT Employment Measures in Productivity Regressions

Model	(1)	(2)	(3)	(4)	(5)
	OLS	<i>Gross output</i> OLS	Trans-log OLS	Exclude finance OLS	FE
Non-IT employment	0.631 (0.011)**	0.523 (0.013)**		0.638 (0.012)**	0.798 (0.008)**
Capital	0.274 (0.009)**	0.340 (0.010)**		0.261 (0.009)**	0.126 (0.006)**
IT employment	0.086 (0.007)**	0.076 (0.008)**	0.087 ^a	0.091 (0.007)**	0.033 (0.004)**
<i>N</i>	36,305	36,305	36,305	33,960	36,305

Notes. Dependent variable is *value added* in all models, except for (2) where we use *gross output*. Each regression also includes *capital*, *non-IT employment*, *industry*, and *year*. Industry controls are at the one-digit SIC level.

(1) is the baseline model. (2) uses *gross output* instead of *value added* as the dependent variable. In (3), we report results using a trans-log rather than a Cobb–Douglas model. In (4), we exclude the financial sector from our sample. In (5), we report results from a fixed effects model.

^aIT output elasticity computed from trans-log estimates; e.g., see Brynjolfsson and Hitt (1995).

** $p < 0.05$.

4.1. Baseline Productivity Estimates Through 2006

We first report baseline estimates using our full data sample from 1987 to 2006. The cross-sectional estimates in Column (1) of Table 10 indicate an elasticity of 0.086 ($t = 12.3$). In Column (2) we show estimates using *gross output* rather than *value added* as the dependent variable. In (3), we use a trans-log rather than a Cobb–Douglas production function. The results from these two columns suggest that our estimates are not overly sensitive to choice of dependent variable or functional form. In Column (4), we remove the finance industry, which is known to be problematic in production function studies due to output measurement concerns and has been omitted in some prior IT-productivity studies for that reason (e.g., Brynjolfsson and Hitt 1996). Given that we have similar results with this sector omitted, our results are not sensitive to the inclusion of this sector.

The estimate from a fixed effects (FE) model, reported in Column (5) of Table 10, is about 0.033 ($t = 8.3$). This is consistent with prior work that suggested a significant component of IT-related productivity contribution is slow changing IT-related organizational practices (Bresnahan et al. 2002) and that much of the “excess return” attributed to IT is due to IT-related intangible assets (Brynjolfsson et al. 2002). These fixed effects results complement rather than replace the OLS results because they are based on different underlying assumptions about what is being measured. The OLS results are comparable to standard growth accounting approaches and allow for the IT coefficient to absorb some of the effects of IT-related organization complements. However, they are subject to biases from unobserved heterogeneity on other dimensions. The fixed effects analyses (as well as the differences analyses presented later) discard this component of IT returns as well as any benefits of IT that are persistent at the firm level over time. They are therefore more conservative econometrically but also likely to substantially underestimate actual IT returns.

In Table 11, we show the results of a total factor productivity (TFP) decomposition, a growth analysis that

has been used in earlier IT productivity research to measure the contribution of different inputs to growth (e.g., see Oliner and Sichel 2000, Brynjolfsson and Hitt 2003). The TFP decomposition numbers are computed using the estimated elasticities from our baseline cross-sectional regression in Column (1) of Table 8 and then multiplying them by the change in input quantity for each input. Total factor productivity is the change in output remaining after the contribution of each factor input has been removed. These results show that the total output in our sample has increased, on average, by 0.25% per year because of increases in IT labor. About half of this rise is due to the change in the input quantity of labor, with the remainder representing increases in multifactor productivity as a result of the excess returns on IT labor. Given that total change in multifactor productivity in our sample is about 1% per year, this indicates that IT labor alone is responsible for 22% of the improvements in TFP observed over our time period. In the next section, we turn toward exploring the effects of endogeneity on our IT estimates.

4.2. Endogeneity

Baseline estimates from Column (1) of Table 8 are reproduced in Column (1) of Table 12. In Columns (2) and (3), we examine the extent to which these results are influenced by reverse causality. Column (2) shows estimates from the Levinsohn–Petrin (LP) estimator, a generalized method of moments (GMM)-based estimator that uses material inputs to control for the effects of unobserved productivity shocks (Levinsohn and Petrin 2003).¹⁸ The Levinsohn–Petrin estimate, 0.077 ($t = 11.0$), is slightly lower than the unadjusted cross-sectional estimate, confirming that the endogeneity of IT hiring imposes an upward bias on unadjusted cross-sectional estimates of IT productivity. However, this analysis confirms two important properties of our measure. First, the productivity contribution of IT labor is still positive and significant, consistent with prior work. Second, despite IT employment likely being subject to a greater endogeneity problem than capital because firms may be able to adjust their employment levels more readily than their capital stock, the effects of endogeneity are small. Indeed, if IT employment is more subject to endogeneity than IT capital, this result implies the

Table 11 Contributions to Growth in Value Added, 1987–2006

Productivity	Estimated contribution ^a
Growth rate of <i>value added</i>	3.38
Contributions from:	
<i>Capital</i>	0.77
<i>IT employment</i>	0.25
<i>Other employment</i>	1.45
<i>Multifactor productivity</i>	0.91

^aAverage annual log difference shown multiplied by 100.

¹⁸ The LP estimator assumes that unobserved productivity shocks affect all variable inputs. The estimation involves a two-stage procedure where changes in materials are used to approximate the unobserved productivity shock, which is then used as a control variable in a productivity estimation. See Olley and Pakes (1996) and Levinsohn and Petrin (2003) for a detailed discussion of this general approach. Estimates were performed using the LEVPET package for STATA.

Table 12 Endogeneity

DV: <i>Value added</i>	Baseline OLS (1)	Levinsohn–Petrin GMM-LP (2)	Arellano–Bond GMM-AB (3)	IT labor quality OLS (4)	IT outsourcing OLS (5)	IT outsourcing OLS (6)
<i>IT employment</i>	0.086 (0.007)***	0.077 (0.007)**	0.041 (0.005)**	0.078 (0.007)***	0.124 (0.039)***	0.122 (0.039)***
<i>IT education</i>				0.010 (0.014)		
<i>IT experience</i>				−0.005 (0.013)		
<i>IT experience</i> ²				−0.005 (0.013)		
<i>Workforce educ</i>				0.192 (0.031)***		
<i>Workforce exp</i>				1.07 (0.189)***		
<i>Workforce exp</i> ²				−0.248 (0.041)***		
<i>IT outsourcing</i>						0.047 (0.024)*
Controls	<i>Industry</i> <i>Year</i>	<i>Industry</i> <i>Year</i>	<i>Industry</i> <i>Year</i>	<i>Industry</i> <i>Year</i>	<i>Industry</i> <i>Year</i>	<i>Industry</i> <i>Year</i>
<i>N</i>	36,273	36,305	25,626	33,889	2,217	2,217
<i>R</i> ²	0.92			0.92	0.85	0.85

Notes. Dependent variable in all regressions is *value added*. Regression also includes *capital*, *non-IT employment*, *one-digit SIC industry*, and *year*.

** $p < 0.05$, *** $p < 0.01$.

endogeneity bias in prior studies of IT value using the productivity framework may be negligible.

The length of our panel also allows us to report Arellano–Bond System GMM estimates in Column (3), a dynamic panel estimator that uses lagged differences as instruments to account for endogenous regressors and was developed specifically with microproductivity measurement in mind (see Blundell and Bond 2000). The coefficient estimate on IT labor from the Arellano–Bond estimator, shown in Column (3) of Table 10, is 0.041 ($t = 8.2$), slightly larger than the fixed effects estimate but lower than the OLS estimates. This appears reasonable because the Arellano–Bond estimator is essentially an optimally weighted average of the regression in levels and the first difference regression (similar to fixed effects) simultaneously estimated by instrumental variables, so we would expect the estimates to fall between the fixed effects and OLS results.

Our data are also uniquely positioned to address other endogeneity concerns present in the IT value literature. In Column (4), we examine whether differences in IT skill composition explain IT returns. We include controls for the education and experience of the IT workforce as well as the education and experience of the firm’s total workforce. Education is computed from the self-reported education levels of the IT workers in our sample. Experience levels for an IT worker in a particular year are estimated

based on adjusting backwards self-reported IT experience levels from 2006, assuming full employment during the interim period. Similar education and experience measures were developed for the firm’s total workforce (IT and non-IT). Because of the scarcity of firm-level human capital data, these data represent a rare opportunity to include both IT and workforce measures in a single regression, allowing us to test whether the IT coefficient reflects higher education levels or other similar attributes. The estimates on IT education and experience are not significantly different from zero. The estimates on workforce education and experience are significant and positive, and the estimate on the experience square term is negative, suggesting that productivity is increasing with experience, but at a diminishing rate. Notably, however, including these measures somewhat lowers the estimated IT elasticity, but by a relatively small amount, from 0.086 to 0.078. Therefore, higher productivity levels associated with IT investment do not appear to be reflecting differences in IT labor quality.

Finally, in Columns (5) and (6), we test how IT outsourcing affects our IT productivity estimates. One concern is that the marginal product of IT would be overestimated if firms outsourced their IT functions—they would receive the full output benefit of IT investment, but the resources producing this output would not appear as IT expenditure, thus appearing to create excess marginal product. Although these analyses

Downloaded from informs.org by [165.123.34.86] on 09 January 2015, at 16:13. For personal use only, all rights reserved.

have been done before, the limited size of IT panels meant that there was often insufficient sample size to get reliable estimates once outsourcing data were matched to IT data. We incorporate recent survey data from more than 200 firms indicating the percentage of their IT budgets dedicated to IT outsourcing. In Column (5), we report baseline estimates from the smaller sample of firms for which IT outsourcing data are available. Firm size from the outsourcing sample is considerably larger than the average firm in our sample, reflected in the higher estimated output elasticity of 0.122 ($t = 3.18$). It is noteworthy, however, that this estimate changes little when we include IT outsourcing levels into our production function, so persistent excess returns to IT spending do not appear to be caused by IT outsourcing.¹⁹


4.3. Time, Size, and Industry Subsamples

In Figure 3, we report results from tests of whether IT returns have diminished after firms completed large waves of investment in business process reengineering in the 1980s and 1990s. We report results from three subsamples: 1987 to 1994, corresponding to the most accurate years of the CITDB and the basis of prior results (e.g., Brynjolfsson and Hitt 2003); 1995 to 1999, the years corresponding to the height of the dot-com boom; and 2000 to 2006. Breaking the sample into these periods allows for direct comparability with earlier results where available and isolates the effects that the dot-com boom may have on our estimates. In Column (1), we show that the elasticity from 1987 to 1994 using our data is 0.053. Prior estimates for this period are in the vicinity of 0.04, and the differences between our estimates and prior estimates can almost entirely be explained by lower measurement error in our data. These elasticity estimates steadily rise (Columns (2) and (3)), reflecting greater IT investment, and double in size in the 2000–2006 period. Moreover, the associated marginal product also appears to be rising in successive time periods, suggesting that not only have investment levels in IT been rising but also that each dollar of IT spending contributes more to overall productivity. Thus, we find no evidence that information technology has contributed less over the post-dot-com era than it had previously, at least for publicly traded firms.

We also compare IT returns between the very large firms that have formed the basis of most prior firm-level IT productivity research and smaller firms.

¹⁹ This analysis also suggests there may be a positive productivity benefit of outsourcing, which would be consistent with firms' realizing the same level of IT services quality at lower cost. However, further analysis and data may be needed to establish this definitively because outsourcing firms are larger than and may differ in other ways from nonoutsourcing firms, and the coefficient on outsourcing is only marginally significant.

Figure 3 Gross Marginal Product of Computer Personnel by Time Period



	(1)	(2)	(3)
Time period	1987–1994	1995–1999	2000–2006
IT employment	0.053	0.073	0.110
Std. error	(0.012)**	(0.010)**	(0.009)**
Marginal product	195.6	273.0	442.7
N	8,547	10,523	17,235


Notes. In thousands of dollars of *value added* per IT employee. Robust standard errors shown in parentheses. Regressions are from baseline model in Column (1) of Table 10 and include controls for *year* and *one-digit SIC industry*. The hypothesis that OLS estimates across time periods are equal is rejected at the $p < 0.001$ level ($F(2, 4,730) = 14.9$). Level estimates are shown by time period.

** $p < 0.05$.

In Figure 4, we compare the contributions of computer capital to the productivity of Fortune 500 firms and non-Fortune 500 firms.²⁰ Although all of our firms are fairly large, there is a considerable difference in size between the two subsamples: our Fortune 500 firms have an average *value added* of \$2.4 billion, but for our non-Fortune 500 firms, the comparable figure is \$526 million. Our estimates suggest that gross marginal product is significantly greater for Fortune 500 firms than for smaller firms. This difference in elasticity is significant ($t = 4.3$, $p < 0.001$). This difference is almost entirely due to a larger elasticity of IT labor in large firms rather than differences in investment levels (which are about the same, relative to size). Fortune 500 firms appear to receive more than double the benefit from the marginal IT dollar. Our results suggest that although IT is productive for both large and moderate size firms, returns to IT may vary substantially across firms of different sizes. There may be a number of explanations for

²⁰ Data on Fortune 500 membership are taken from the Compustat database. We identify a firm as belonging to the Fortune 500 if the firm appears in the Fortune 500 in *any year* between 1987 and 2006. Thus, our Fortune 500 identifier is meant primarily to distinguish a firm type as opposed to literally being in the Fortune 500 in an arbitrarily chosen year.

Figure 4 Gross Marginal Product of Computer Personnel by Firm Size



	F500 (1) OLS	Non-F500 (2) OLS	F500 (3) FE	Non-F500 (4) FE
IT employment	0.126	0.077	0.022	0.033
Std. error	(0.017)**	(0.008)**	(0.006)**	(0.005)**
Marginal product	530.5	258.9		
N	10,533	25,772	10,533	25,772

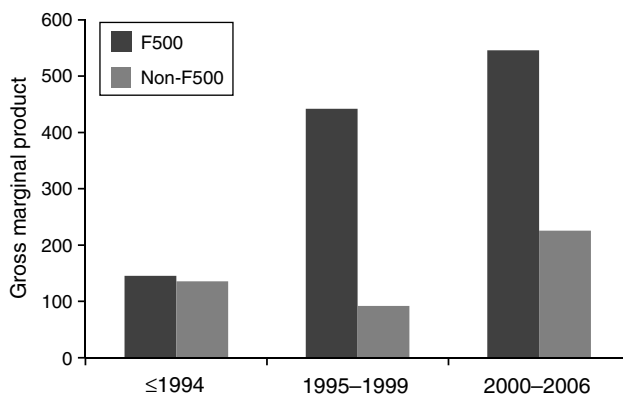
Notes. In thousands of dollars of *value added* per IT employee. Regressions are from baseline model in Column (1) of Table 10 and include controls for *year* and *one-digit SIC industry*. The hypothesis that the OLS coefficients are equal across samples in (1) and (2) is rejected at the $p < 0.001$ level ($t = 4.31$).

** $p < 0.05$.

this finding, including that enterprise systems (e.g., SAP) are disproportionately employed in larger firms and have significant effects on productivity (Hitt et al. 2002) or simpler arguments such as economies of scale in IT purchasing or deployment. Our estimates in Columns (3) and (4), however, provide some evidence for the first interpretation because the reduction in the size of the large firms' estimate after including fixed effects suggests that much of the difference in coefficient estimates between the two subsamples may be associated with slow changing IT-related organizational practices.

To shed some additional light on the nature of these differences, we further break down these results by time period in Figure 5. Both the marginal product and elasticity of IT investment for midsize and large firms are comparable between 1987 and 1994. However, these values begin to diverge in subsequent years, with the elasticity and marginal product of IT becoming much greater for large firms after 1995 and through 2006. In the later time period, the marginal product of IT investment in large firms is

Figure 5 Changes over Time in Gross Marginal Product of Computer Personnel in Firms of Different Sizes



Sample	F500 (1)	Non-F500 (2)	F500 (3)	Non-F500 (4)	F500 (5)	Non-F500 (6)
Years	1987–1994	1987–1994	1995–1999	1995–1999	2000–2006	2000–2006
OLS IT coefficient	0.037	0.044	0.103	0.031	0.126	0.062
Std. error	(0.015)**	(0.013)***	(0.020)***	(0.010)***	(0.020)***	(0.016)***
Marginal product	145.7	135.9	442.0	92.0	546.1	225.6
FE IT coefficient	0.013	0.042	0.027	0.038	0.049	0.043
Std. error	(0.008)*	(0.011)***	(0.013)**	(0.010)***	(0.012)***	(0.010)***
N	3,650	4,897	2,859	7,664	4,024	13,211

Notes. In thousands of dollars of *value added* per IT employee. Reported OLS value is coefficient estimate on $\log(\text{IT employment})$ in OLS productivity regressions that also include logged values of *capital*, *non-IT employment*, *one-digit SIC industry*, and *year* as independent variables. Reported FE value is coefficient estimate on $\log(\text{IT employment})$ in FE productivity regressions that also include logged values of *capital*, *non-IT employment*, *one-digit SIC industry*, and *year* as independent variables.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

more than double that of smaller firms. However, as with the above pattern of estimates, these differences disappear after including fixed effects, suggesting that much of this IT-related value is associated with time-invariant firm attributes that are correlated with IT inputs, such as investments in IT-related intangible assets. The observed pattern of estimates also suggests that IT returns may be slower to materialize in large firms if large firms require more time to make complementary organizational investments. We explicitly test this in the next section.

Given that we have a very long (20-year) panel, fixed effects estimates may not provide a suitable control for idiosyncratic firm characteristics that may evolve over time. In Table 13, therefore, we report a number of estimates using one-period and longer differences. In addition to providing the same control for firm-specific effects as a fixed effects analysis, the coefficient estimates at varying difference lengths can be interpreted as a comparison of the short run (first differences) versus long run (three-year or more differences; see Bartelsman et al. 1994). Longer differences may also be less subject to biased estimates from measurement error (Greene 1993). Columns (1) and (2) show the differenced estimates in the CITDB sample and in Fortune 500 firms. The results are similar to those from earlier research using comparable samples (Brynjolfsson and Hitt 2003)—the output contribution of IT appears to be greater over longer difference lengths, which is consistent with the idea that firms require time to install organizational practices that complement IT investments. In Column (3), we show results for non-Fortune 500 firms, and in Column (4), we show differences

results from the complete sample of firms. Unlike the results in Columns (1) and (2), the output contribution of IT does not grow over time in samples that include smaller firms. Rather, IT returns are significant and positive using this difference specification for all firms, but they appear to be slightly higher over one-year differences than for seven-year differences. The absence of growth in output contribution over long differences is consistent with the interpretation that IT adjustment costs are smaller in smaller firms. A comparison of the long and short difference results also supports our prior assertion that the coefficients are not appreciably biased by measurement error. If measurement error were substantial, the long difference results would be considerably larger than the short difference results. This provides additional evidence that our *IT labor* variable has desirable measurement properties.

In Columns (5) and (6), we separate the results using Fortune 500 firms into seven-year periods of 1991–1998 and 1999–2006. The results from the earlier time period show a more consistent rise in output contribution than the results in Column (1) and are close to comparable estimates in earlier research using CITDB data (Brynjolfsson and Hitt 2003). The evidence in the latter time period shows more consistent returns that do not appear to grow over time, suggesting that the timing delays because of the development of organizational complements may play a less important role than they did in earlier decades.


In Figure 6, we show marginal product by one-digit North American Industry Classification System (NAICS) economic sector. The one-digit sector level conceals a great deal of heterogeneity at finer

Table 13 Estimates with Varying Difference Lengths

	(1)	(2)	(3)	(4)	(5)	(6)
DV: <i>Value added</i>	CITDB	F500 firms	Smaller firms	All firms	F500 ≤1998	F500 >1998
1 Year differences	0.027 (0.009)	0.017 (0.006)	0.052 (0.007)	0.044 (0.006)	0.016 (0.006)	0.020 (0.012)
2 Year differences	0.020 (0.009)	0.020 (0.007)	0.040 (0.008)	0.035 (0.006)	0.012 (0.008)	0.035 (0.012)
3 Year differences	0.021 (0.009)	0.025 (0.008)	0.037 (0.008)	0.033 (0.007)	0.017 (0.009)	0.039 (0.013)
4 Year differences	0.027 (0.010)	0.028 (0.008)	0.038 (0.009)	0.035 (0.007)	0.015 (0.010)	0.047 (0.014)
5 Year differences	0.030 (0.012)	0.023 (0.009)	0.033 (0.009)	0.030 (0.007)	0.018 (0.011)	0.035 (0.014)
6 Year differences	0.033 (0.013)	0.024 (0.010)	0.039 (0.010)	0.035 (0.007)	0.023 (0.012)	0.032 (0.014)
7 Year differences	0.035 (0.015)	0.028 (0.011)	0.038 (0.011)	0.036 (0.008)	0.024 (0.013)	0.035 (0.015)

Notes. Each cell corresponds to a separate productivity regression. The reported value is the coefficient estimate on $\log(\text{IT employment})$ in productivity regressions that also include logged values of *capital*, *non-IT employment*, *one-digit SIC industry*, and *year* as independent variables.

Figure 6 Gross Marginal Product of Computer Personnel by Sector



Sector	IT coefficient	Std. error	Marginal product	N
Construction	-0.165	0.200	-656.1	37
Nondurable manufacturing	0.083	0.011**	314.4	11,142
Durable manufacturing	0.032	0.015**	173.7	6,383
Transportation and utilities	0.006	0.014	29.0	3,458
Wholesale trade	0.059	0.041	196.4	1,506
Retail trade	0.067	0.023**	222.6	3,670
Financial services	-0.006	0.039	57.3	2,345
Nonfinancial services	0.151	0.015**	266.3	7,764

Notes. In thousands of dollars of *value added* per IT employee. Robust standard errors shown. The hypothesis that the coefficient estimates across sectors are equal is rejected at the $p < 0.001$ level ($F(74, 730) = 4.91$) Regressions are from baseline model in Column (1) of Table 10 and include controls for *year* and *one-digit SIC industry*.

** $p < 0.05$.

industry levels but is chosen to facilitate direct comparisons with prior work. Brynjolfsson and Hitt (1996) present some evidence suggesting higher returns in computer-producing industries and in firms outside the service sector but do not have the sample size required to produce significant results when the data is subdivided by industry. Our estimates are significant and provide evidence that information technology spending has had a greater impact in manufacturing sectors than in service sectors and that within manufacturing, computer capital produces the most value added in nondurable manufacturing.²¹ This is inconsistent with a disproportionate return in computer-producing firms, which are concentrated in durable manufacturing. The chart does not show estimated returns in mining/construction and finance, although we report them in the accompanying table. The estimation difficulties in finance and mining are consistent with earlier work (Brynjolfsson and Hitt 1996) and are potentially the result of output measurement problems in finance and small sample sizes in mining.

²¹ IT returns are significantly higher for computer-producing industries in the broader sample, but we cannot reject the hypothesis that returns for computer-producing industries within manufacturing are equal to IT returns for other manufacturing industries.

5. Discussion

The importance of collecting new data for measuring the digital economy has been identified as a key theme for advancing research on technology and productivity (Haltiwanger and Jarmin 2000, Mesenbourg and Atrostic 2000, Atrostic et al. 2001). In this analysis, we report a number of new findings about IT productivity enabled by the development of new, firm-level measures of IT expenditures based on a large IT employment sample. These measures have a number of desirable properties, including the availability of large samples both in cross section and in time series, relatively low measurement error, comparability with prior measures, and their extension into recent time periods.

First, the length of the panel allows us to use new estimators from the microproductivity literature to establish that IT productivity results appear to suffer from a relatively small endogeneity bias. Because of the difficulty in finding clean instruments, the endogeneity question has been persistent in the IT productivity literature, making it difficult to make causal interpretations about existing IT productivity estimates; although new estimators have become available that address these issues, they require longer panels. Our findings from using these estimators suggest that although higher productivity levels are likely

to have a reverse effect on IT spending, this has a limited effect on IT productivity estimates; they also suggest bounds for the effects of endogeneity on IT productivity estimates.

Second, we address some important limitations imposed by sample restrictions in earlier IT data. We showed that the productivity benefits of IT have persisted after the dot-com bust and in fact appear to be higher now than in previous periods, especially in larger firms. Furthermore, because our sample includes mid-sized public firms, we could test the hypothesis that the long-run contribution of IT investment is different in large firms. Namely, an important reason for the rising levels of IT value after 2000 is continued growth in large firms. This finding is consistent with the hypothesis that IT-related intangibles take longer to develop in larger firms and that over the last decade large firms have continued to build valuable IT-related intangibles.

The literature on differences in IT productivity in smaller firms is still developing and deserves further attention.²² For managers in firms of all sizes, these results have implications for the pattern of IT returns that they can expect over time. Higher adjustment costs in large firms suggest that the payoff from enterprise-wide IT investment takes longer to materialize in large firms. Further unpacking the detailed mechanisms behind these size-related differences in IT productivity provides scope for future research. These findings also have both managerial and policy implications for industrial organization. Higher IT returns in large firms signal rising barriers to entry for smaller firms. There is some emerging evidence of this in the literature—using industry-level data, Saunders (2010a) finds that higher IT intensity in an industry is associated with more large-firm expansion relative to small entry and that IT intensive industries tend to be more concentrated in large firms. Greater returns to IT investment for larger firms have implications for managers interested in understanding scale-based sources of advantage in their industries.

Although there are a number of limitations to our data, including possible differences between our observed population and the broader population of IT workers, none of our tests thus far suggests that

these differences are material either for the absolute levels of the measures or how they perform in productivity regressions. There is some indication that the workers we observe are more likely to change jobs than other workers, but this can only introduce bias if this behavior is idiosyncratic to specific IT workers at specific firms. However, as we and other researchers employ these measures, their advantages and limitations in other modeling frameworks may become more apparent.

References

- Akerberg, D., K. Caves, G. Frazer. 2006. Structural identification of production functions. Working paper, University of California, Los Angeles, Los Angeles.
- Applegate, L., J. Cash, D. Mills. 1988. Information technology and tomorrow's manager. *Harvard Bus. Rev.* (November–December) 128–136.
- Aral, S., E. Brynjolfsson, D. Wu. 2006. Which came first, IT or productivity? Virtuous cycle of investment and use in enterprise systems. Working paper, New York University, New York.
- Arellano, M., S. Bond. 1991. Some tests of specification for panel data: Monte Carlo evidence and an application to employment equations. *Rev. Econom. Stud.* 58(2) 277–297.
- Arrow, K. 1962. Economic welfare and the allocation of resources to invention. R. R. Nelson, ed. *The Rate and Direction of Economic Activity: Economic and Social Factors*. Princeton University Press, Princeton, NJ, 609–626.
- Atriotic, B. K., J. Gates, R. Jarmin. 2001. Measuring the electronic economy: Current status and next steps. Working Paper 00–10, U.S. Census Bureau, Washington, DC.
- Bartelsman, E., R. Caballero, R. Lyons. 1994. Customer- and supplier-driven externalities. *Amer. Econom. Rev.* 84(4) 1075–1084.
- Bloom, N., R. Sadun, J. Van Reenen. 2008. American's do I.T. better: US multinationals and the productivity miracle. *Amer. Econom. Rev.* Forthcoming.
- Blundell, R., S. Bond. 2000. GMM estimation with persistent panel data: An application to production functions. *Econometric Rev.* 19(3) 321–340.
- Bresnahan, T., S. Greenstein. 1996. Technical progress and co-invention in computing and the uses of computers. *Brookings Papers Econom. Activity* 1996 1–83.
- Bresnahan, T., E. Brynjolfsson, L. Hitt. 2002. Information technology, workplace organization, and labor demand. *Quart. J. Econom.* 117(1) 339–376.
- Brynjolfsson, E. 1993. The productivity paradox of information technology. *Comm. ACM* 36(12) 67–77.
- Brynjolfsson, E., L. Hitt. 1996. Paradox lost? Firm-level evidence on the returns to information systems spending. *Management Sci.* 42(4) 541–558.
- Brynjolfsson, E., L. Hitt. 2000. Beyond computation: Information technology, organizational transformation and business performance. *J. Econom. Perspect.* 14(4) 23–48.
- Brynjolfsson, E., L. Hitt. 2003. Computing productivity: Firm-level evidence. *Rev. Econom. Statist.* 85(4) 793–808.
- Brynjolfsson, E., A. Saunders. 2010. *Wired for Innovation*. MIT Press, Cambridge, MA.
- Brynjolfsson, E., D. Fitoussi, L. Hitt. 2006. The information technology iceberg. Working paper, MIT, Cambridge, MA.
- Brynjolfsson, E., L. Hitt, S. Yang. 2002. Intangible assets: Computers and organizational capital. *Brookings Papers Econom. Activity* 2002(1) 137–198.

²² Somewhat in contrast to our findings, a recent study finds that the link between computer networks and labor productivity is stronger in smaller manufacturing plants than in larger ones and that IT investment has a negative effect on the productivity of large plants (Hyatt and Nguyen 2010). There are a number of explanations for the differences between these two findings. The higher IT returns we attribute to large firms are not captured until the later part of the panel, which postdates much of the Hyatt and Nguyen sample that is based on 1999 data. Second, a plant-level analysis does not capture some of the IT benefits that would be captured by large firms, such as superior coordination among multiple establishments.

- Cheng, Z., B. R. Nault. 2007. Industry level supplier-driven IT spillovers. *Management Sci.* 53(8) 1199–1216.
- Chwelos, P., R. Ramirez, K. L. Kraemer, N. P. Melville. 2010. Does technological progress alter the nature of information technology as a production input? New evidence and new results. *Inform. Systems Res.* 21(2) 392–408.
- Dedrick, J., V. Gurbaxani, K. Kraemer. 2003. Information technology and economic performance: A critical review of the empirical evidence. *ACM Comput. Surveys* 35(1) 1–28.
- Dewan, S., C. Min. 1997. The substitution of information technology for other factors of production. *Management Sci.* 43(12) 1660–1675.
- Gordon, R. 2000. Does the new economy measure up to great inventions of the past? *J. Econom. Perspect.* 14(4) 49–74.
- Greenan, N., J. Mairesse. 2000. Computers and productivity in france: Some evidence. *Econom. Innovation New Technol.* 9(3) 275–315.
- Greene, W. 1993. *Econometric Analysis*. Macmillan, New York.
- Hall, B. 1990. The manufacturing sector master file: 1959–1987. NBER Working Paper 3366, National Bureau of Economic Research, Cambridge, MA.
- Haltiwanger, J., R. S. Jarmin. 2000. Measuring the digital economy. E. Brynjolfsson, B. Kahin, eds. *Understanding the Digital Economy: Data, Tools, and Research*, MIT Press, Cambridge, MA, 13–33.
- Henderson, R. 1993. Underinvestment and incompetence as responses to radical innovation: Evidence from the photolithographic alignment equipment industry. *RAND J. Econom.* 24(2) 248–270.
- Hitt, L. M., D. Wu, X. Zhou. 2002. Investment in enterprise resource planning: Business impact and productivity measures. *J. Management Inform. Systems.* 19(1) 71–98.
- Hyatt, H., S. Nguyen. 2010. Computer networks and productivity revisited: Does plant size matter? Working paper, Center for Economic Studies, U.S. Census Bureau, Washington, DC.
- InformationWeek. 2006. Multi-year comparisons of informationweek 500 survey data reveal IT trends. (September 12), <http://www.informationweek.com/news/management/showArticle.jhtml?articleID=192600107>.
- Ito, H. 1995. The structure of adjustment costs in mainframe computer investment. Working paper, Stanford University, Stanford, CA.
- Jorgenson, D., M. Ho, K. Stiroh. 2008. A retrospective look at the U.S. productivity growth resurgence. *J. Econom. Perspect.* 22(1) 3–24.
- Lee, B., A. Barua, A. Whinston. 1997. Discovery and representation of causal relationships in MIS research: A methodological framework. *MIS Quart.* 21(1) 109–136.
- Levinsohn, J., A. Petrin. 2003. Estimating production functions using inputs to control for unobservables. *Rev. Econom. Stud.* 70(2) 317–341.
- Lichtenberg, F. 1995. The output contributions of computer equipment and personnel: A firm-level analysis. *Econom. Innovation New Technol* 3(3–4) 201–218.
- Mata, F., W. Fuerst, J. Barney. 1995. Information technology and sustained competitive advantage: A resource based analysis. *MIS Quart.* 19(4) 487–505.
- McElheran, K. 2010. Incumbent response to process innovation: The case(s) of e-business adoption. Working paper, Harvard Business School, Boston.
- Melville, N., K. Kraemer, V. Gurbaxani. 2004. Information technology and organizational performance. *MIS Quart.* 28(2) 283–322.
- Mesenbourg, T., B. K. Atrostic. 2000. Measuring the U.S. digital economy: Theory and practice. Working paper, U.S. Census Bureau, Washington, DC.
- Oliner, S., D. Sichel. 2000. The resurgence of growth in the late 1990's: Is information technology the story? *J. Econom. Perspect.* 14(4) 3–22.
- Olley, S., A. Pakes. 1996. The dynamics of productivity in the telecommunications equipment industry. *Econometrica* 64(6) 1263–1298.
- Roach, S. 1991. Services under siege—The restructuring imperative. *Harvard Bus. Rev.* (September–October) 82–92.
- Saunders, A. 2010a. Has information technology leveled the competitive playing field? Working paper, Wharton School, University of Pennsylvania, Philadelphia.
- Saunders, A. 2010b. Valuing IT-related intangible capital. Working paper, Wharton School, University of Pennsylvania, Philadelphia.
- Steiner, T. D., D. B. Teixeira. 1990. *Technology in Banking: Creating Value and Destroying Profits*. McGraw-Hill, New York.
- Stiroh, K. 2005. Reassessing the impact of IT in the production function: A meta-analysis and sensitivity tests. *Ann. Econom. Statist.* (79/80) 529–561.
- Stiroh, K. 2008. Information technology and productivity: Old answers and new questions. *CESifo Econom. Stud.* 54(3) 358–385.
- Tambe, P., L. M. Hitt. 2012. Job hopping, information technology spillovers, and productivity growth. Working paper.
- Tambe, P., L. M. Hitt, E. Brynjolfsson. 2012. The extroverted firm: How external information practices affect innovation and productivity. *Management Sci.* Forthcoming.
- Tushman, M., P. Anderson. 1986. Technological discontinuities and organizational environments. *Admin. Sci. Quart.* 31(3) 439–465.
- Varian, H. 1992. *Microeconomic Analysis*, 3rd ed. W. W. Norton, New York.
- Wooldridge, J. 2002. *Econometric Analysis of Cross-Section and Panel Data*. MIT Press, Cambridge, MA.