

Assessing the Rate of Change in the Enterprise Database System Market Over Time Using DEA

Timothy R. Anderson¹, Keith Hollingsworth², Lane Inman^{1,3}

¹ Engineering and Technology Management Department, Portland State University, Portland, OR 97207-0751, USA

² Economics and Business Administration Department, Morehouse College, Atlanta, GA 30314, USA

³ VERITAS Software Incorporated, Mountain View, CA 94043, USA

Abstract-- This paper builds upon previous work of using Data Envelopment Analysis (DEA) to measure incremental innovation in technology by applying it to the Online Transaction Processing Market. A variable returns to scale DEA model is utilized to determine an annual rate of change in benchmarks based on data provided by the Transaction Processing Performance Council. This rate of change may then be used to forecast possible future performance trendsetters of the TPC.

I. INTRODUCTION

In Anderson, et al, [2], Data Envelopment Analysis (DEA) is used to estimate the annual rate of change of new microprocessor performance and illustrates consistency with Moore's Law, a widely held rule stating a steady rate of microprocessor attribute change over time. This paper implements similar techniques to measure the rate of change of on-line transaction processing (OLTP) database systems. Transaction Processing Performance Council (TPC) system benchmarks are used to illustrate variable returns to scale (VRS), DEA and the technology rate of change approach. This technique is readily applied to more complex applications.

A. Technology Forecasting

Technological forecasting provides procedures for data collection and analysis to predict future technological developments and the impacts such developments will have on the environment and lifestyles of mankind [21]. The measurement of technology in order to identify a rate of change has received much research. Many different approaches currently attempt to address this challenge. All of which require a set of assumptions and a priori measures.

Before technological progress can be measured, those factors that determine progress must be accurately assessed. This lends itself to a large amount of subjective interpretation, which introduces chances of imprecision. There have been a number of works that have implemented forecasting methodologies in conjunction with factor analysis and expert opinion to come up with these factors [9, 15, 17, 18].

A major assumption of many quantitative technology forecasting techniques is that technology trends are a valid model for future trends. Although often a reasonable assumption, it is a limitation of the models no less. With the

advent of sudden technological changes in an industry, much can be seen in the ways of differing results. This is particularly true in Information Technologies and other rapid technology advancing fields [9].

The level of system aggregation is also a key aspect to consider when performing technology measurement. Whether or not it is desirable to consider components of the overall system or the system as a whole [15]. Additional considerations are considerations of major innovations or merely gradual improvements over time. These gradual improvements typically denote a continuous process of improvement of technologies.

Popular methods of forecasting include multiple regression analysis to determine a function of various factors to predict the rate of change of technology and target reasonable future elements, based on a given dataset [9, 13, 15]. Much of this can be derived from the concept of a technology surface representing the state of the art, SOA [1, 8]. These models are limited by fixed weighting imposed by the multiple regression technique. Other issues include the SOA surface does not change, the rate of technology change is assumed to be constant, historical continuity prevails, and trade-offs in the design process are not explicitly considered. Much work has been put towards making the elements dimensionless as well as utilization of various forms of factor analysis to more accurately represent points of technology.

The Hedonic approach operates on the basic hypothesis that a product or services utility lies in its essential attributes. Difference in commodity pricing is related to differences in their quality characteristics. This is then decomposed into "quality" or "technological change" and changes can be estimated through regression over time. However it suffers the limitation of assuming that all user utility functions are equal at given values [15].

Additional work has been conducted to address some of these issues in a composite approach [16-18]. It builds upon previous work and splits it into two approaches. The first approach is a holistic index of technology that is a yardstick for advances of the technology surface structure over time. Also, the holistic index of technology can indicate deep structural changes of technical knowledge over time. In short one approach effectively seeks to measure the advances in the applications of known physical laws while the other seeks to map the further understanding of those laws. These surfaces

are determined through surfaces with constant probability density given performance and design variables. Progress is then related through Mahalanobis distances.

Although the above-mentioned measures have brought many insights to technology forecasting, they often suffer from being based upon the requirement of a fixed weighting in the determination of linear functions. DEA offers a means to allow for flexible weighting as determined by the data set and allows for changing weights over time. For the purpose of this paper, the efficiency frontier of DEA is considered to be synonymous with a technology frontier. This frontier is a piecewise-linear, convex combination of actual observations requiring a minimal set of assumptions.

B. Description of Dataset

As database systems have become more prevalent, measuring their performance has become increasingly important. In the early 1980's, automated teller machine (ATM) networks marked one of the first widespread moves in end-user business transaction automation through information systems. This rapidly expanded to point of sale (POS) devices and registers in gas stations, retail outlets, and grocery stores. As time progressed, transactions were brought closer to the consumer through the Internet resulting in the explosion of e-commerce. Contrary to their batch-processing brethren of the 1960's and 1970's, OLTP database systems involved live transaction processing. To provide accurate and consistent benchmarks for the wide variety of available OLTP database systems, the TPC was formed in the early eighties. Their role eventually evolved to that of a governing body of benchmark administration to further ensure accuracy and integrity.

The TPC has created a number of benchmarks that have evolved over time to provide better indicators of performance. Currently, the most common measure of performance for OLTP database systems is TPC-C, which utilizes a combination of five concurrent transactions centered on order-entry activities. Designed to test an array of system components, the tests are characterized by parallel execution of multiple transactions types, multiple on-line sessions, significant disk input and output, on-line deferred transaction execution modes, transaction properties, non-uniform distribution of data, databases consisting of multiple tables of varying sizes, attributes, and relationships, and data access and update contention. The total costs of ownership are also considered within the benchmark including purchase price for the solution and service costs [20].

The TPC-C benchmarks provide two measures: overall system cost and transactions-per-minute-C, tpmC. Once a vendor submits their benchmarks to the organization they are made publicly available on the TPC's website: www.tpc.org. The dataset in this paper was collected November 20, 2000.

II. METHODOLOGY

A. Introduction to DEA

During the past two decades, over 1500 papers have been published using DEA [19]. It was initially developed in 1978 as a means of measuring productivity in a multi-input/multi-output environment [6]. The Charnes-Cooper-Rhodes, CCR, model creates a data "envelope" around a dataset that serves as an efficiency frontier. Each observation, or decision-making unit (DMU), resting on the envelope's frontier is considered efficient. Efficiencies of those DMUs not resting on the frontier are determined through the calculation of lessened input or additional output required to relocate the DMU to the frontier established by the efficient DMUs. This can be illustrated with the dataset provided in Table 1.

The data of Table I is illustrated in Figure 1. In this example the efficiency frontier is determined by DMUs A, B, C, and D assuming convexity and variable returns to scale (VRS) [4]. Convexity dictates that if C and D exist then it is possible for any linear combination of C and D to exist. VRS indicates that it is possible to have differing returns to scale on the efficiency frontier. In this example, a constant returns to scale model would result in a ray extending from the origin through B as the frontier.

TABLE I
SAMPLE DMU SUBSET

Obs	Name	Avail. Year	Total Cost (\$)	tpmC
A	Compaq ProLiant 3000-6/600-1P	1999	160643	8050
B	Compaq ProLiant ML570	2000	201717	20207
C	Dell PowerEdge 6450	2000	334936	31231
D	Unisys e-@ction Enter.	2000	797935	61390
E	Compaq ProLiant 3000 6/450-512 1	1998	176042	6290
F	Unisys Aquanta QR/2V Server	1998	424297	19118




Figure 1 Example DMU Set Plotted

The maximization of output for a given input is termed output orientation. To find the efficiency for E, a virtual observation E' is mapped on the frontier consisting of a linear combination of 37% B and 63% A that results in an efficient output of 12548. Compared to E's actual output of 6290, E

should have produced 2.00 times as much tpmC for the money spent. The efficiency of E is determined to be 1/2.00 or 0.50. The efficiency of F is determined in the same manner using a combination of C and D. Evaluation of F' reveals that F should create 1.94 times as much tpmC for the money spent resulting in an efficiency of 0.52.

This single-input, single-output application could be analyzed graphically but most applications would require more inputs and/or outputs that necessitate a more sophisticated approach. In practice, the linear combinations used to calculate the efficiency scores are found by performing the basic output-oriented linear programming formulation for each DMU.

$$\begin{aligned}
 & \max \phi, \\
 & \text{s.t.} \quad \sum_{j=1}^n x_{i,j} \lambda_j \leq x_{i,0}, \quad \forall i \in \{1, \dots, m\} \\
 & \quad \sum_{j=1}^n y_{r,j} \lambda_j \geq \phi y_{r,0}, \quad \forall r \in \{1, \dots, s\} \\
 & \quad \sum_{j=1}^n \lambda_j = 1, \\
 & \quad \lambda \geq 0.
 \end{aligned} \tag{1}$$

In this case, $x_{i,j}$ refers to the i 'th input of system j and $y_{r,j}$ corresponds to the r 'th output of system j . The currently evaluated DMU's (or TPC submission's) inputs and outputs are denoted by $x_{i,0}$ and $y_{r,0}$. The variable, λ_j , indicates the amount of DMU j used in setting a performance target for the currently evaluated system. The scalar quantity, ϕ , corresponds to the efficiency where $\phi=1$ indicates that the DMU is radially efficient and $\phi>1$ indicates that the DMU should be achieving more output.¹ The final constraint, $\sum_{j=1}^n \lambda_j = 1$, enforces a variable return to scale. For a more comprehensive treatment of DEA, the interested reader is referred to [7, 10, 11].

B. Technical Change Model

The technical change model is presented by Anderson, et al. and is based on the movement of the efficiency frontier over time [2]. An examination of Figure 2 illustrates the movement of this frontier over time between 1998 and 2000. Here it is easy to see that those DMUs (E and F) considered efficient in the year 1998 are no longer deemed efficient in the year 2000.²

¹ The input-oriented DEA model with an efficiency score, θ , ranging between 0 and 1.0 is probably more commonly used. In the Constant Returns to Scale case, $\phi=1/\theta$. This relationship does not precisely hold in the case of Variable Returns to Scale.

² The sample data set is extracted from the original data set. Only TPC submissions that were efficient at the time of submission were used so E and F were efficient relative to preceding submissions.




Figure 2 - Example Movement of Efficiency Frontier

As had been previously described, the expected output for the given inputs of system F correspond to outputs which should be approximately twice as much as what was achieved by F ; therefore technical progress of database systems at this price point has effectively doubled over the past two years. In other words, if one spent the same amount today, one should expect twice the performance that one received in the past.

The technical change through time can be tracked through an extension of the model proposed by Anderson et al. [2] which relates the radial efficiency (Φ), released radial efficiency (ϕ_0), the coefficient of technical progress (β), and the time between the observation and those which comprise the efficiency frontier (τ).

$$\Phi = \beta^\tau \phi_0 \tag{2}$$

The coefficient of technical progress, β , provides a metric as to how much additional output would be achieved in the future at time τ relative to the point of release. Since β corresponds to the technical progress, β is calculated using only those observations considered efficient at the time of release or, more precisely, when ϕ_0 is equal to one.

For example, system F was on the efficiency frontier in 1998, but in 2000 the frontier has moved. Its virtual counterpart, F' , provides 1.94 times the output, tpmC, for the same cost. The elapsed time between F and the frontier's segment against which F is compared is two years, and thus the calculation of the coefficient of technical progress is calculated below.

$$1.94 = \beta^2 1.0 \tag{3}$$

This leads to a coefficient of technical progress of $\beta = 1.39$. In other words, for a given input, one would expect 39% more output per year after the year that the observation was on the efficiency frontier.³

³ The TPC application examined in this paper is a single-input, single-output DEA model, but DEA and this approach are inherently designed for

System E's evaluation is a little more complicated. Since the efficiency frontier segment E is evaluated against includes systems released in both 1999 and 2000 we take this into account. Equation (4) illustrates the calculation of a weighted average of times to ascertain the effective time since the frontier, τ_E .

$$\tau_E = \frac{\sum_{j=n} \lambda_j (\tau - \tau_j)}{\sum_{j=n} \lambda_j} \quad (4)$$

In VRS, this equation is reduced to Equation 5 since VRS fixes $\sum \lambda$ to 1 resulting in the following.

$$\tau_E = \sum_{j=n} \lambda_j (\tau - \tau_j) \quad (5)$$

Application to system E of (5) yields the following:

$$\tau_E = 0.63 \times (1999 - 1998) + 0.37 \times (2000 - 1998) = 1.37 \quad (6)$$

Thus the weighted time passed relative to the setting of the efficiency frontier, τ_E , is 1.37 years. Once this effective time between the observation and the efficiency frontier have been calculated we can use this in conjunction with (6):

$$2.00 = \beta^{1.37} \times 1.0 \quad (7)$$

This results in a technical progress, β , of 1.65. The resulting coefficient of technical progress is then $\beta = 1.65$.

In the real world, tests that are often relatively close to each other may fall in different years. Such would be the case if one observation takes place in December 1999 and the other takes place in January 2000. To better address this issue, we will use the day as the unit of time denoted by τ . Then, the annual coefficient of technical progress is calculated using (8).

$$\beta_{annual} = \beta_{daily}^{365.24} \quad (8)$$

This results in a more granular and precise tabulation of the rate of technical progress.

C. Results

1) Overview

Overall, 191 TPC-C submissions were considered valid DMUs. As mentioned previously, DMUs that were not considered efficient, those with a radial efficiency (ϕ_0) of over

both multiple inputs and outputs. $\beta = 1.39$ would then be interpreted as indicating that the passing of a year should result in a system achieving at least 39% more of each output (perhaps various performance metrics) using no more of any input (perhaps a variety of cost components).

1.0 on their date of availability, were dropped from further analysis. Of these 191 DMUs, sixty-five defined the efficiency frontier at fifty-five points in time. On occasion, multiple efficient benchmarks were simultaneously made available.

2) Efficient TPC Submissions

Figures 3 and 4 summarize the movement of the efficiency frontier. A point of interest is the presence of a DMU dated March 31, 1997 on the final efficiency frontier. Upon examination of Figure 4, it appears that this point is the anchor for performance. Perhaps this is the minimum "cost of entry" necessary to perform the benchmarks. Also, this low-end system no longer has competition. Since TPC is typically used to indicate how powerful a system solution is vendors may not target this area of the frontier. Additionally, this tpmC benchmark is based on Rev 3.2 as opposed to the later revisions. This could indicate some discrepancies amongst the sub-revisions of the test.

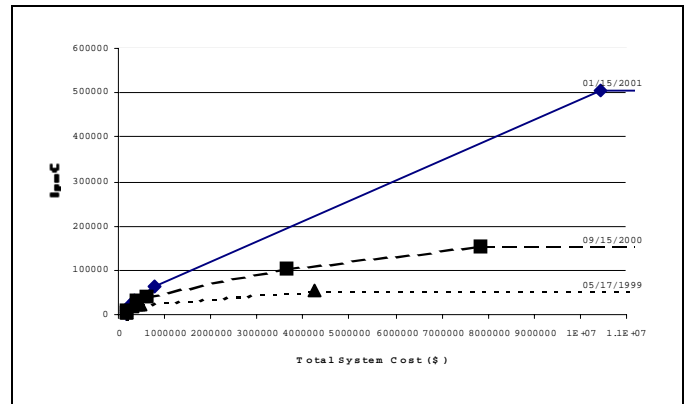


Figure 3 – Movement of Efficiency Frontier Over Time

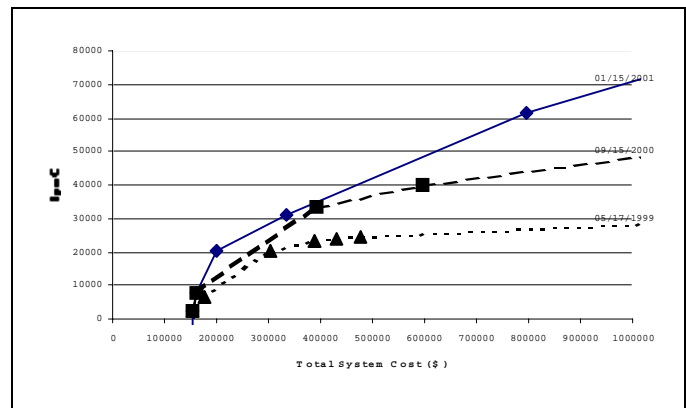


Figure 4 - Closer Examination of Efficiency Frontier

3) Rate of Change Estimates

Each TPC submission analyzed can be used to form a separate estimate of the annual rate of change, β , in the TPC market. Higher values of β correspond to a rapidly advancing

frontier. A high value of β , indicates that the particular TPC submission, while originally efficient relative to all previous submissions, may have had performance that was greatly exceeded in a short period of time.

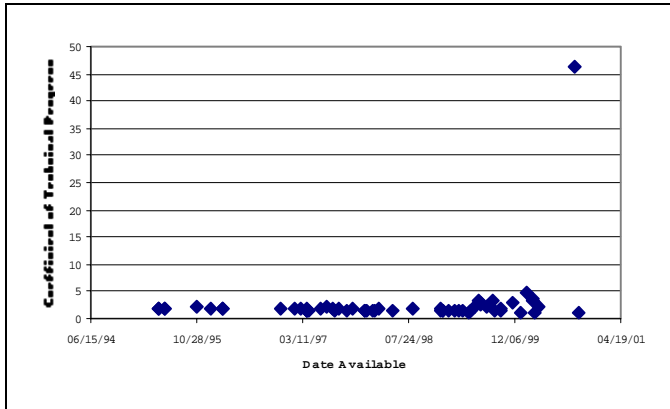


Figure 5 - Coefficient of Technical Progress

Rate of change estimates are illustrated in Figure 5. The annual rate of change estimate for each submission is relatively consistent except for one extreme value dated, June 1, 1999. The high rate of change is due to the fact that this submission's performance was greatly exceeded within a relatively short period of time (9/15/2000). This new system used Microsoft Windows 2000 in a high-end configuration to achieve much higher performance. This combination of higher performance and a short time between the submission results in an extremely high value of a technical progress, β , and could be an indicator of a possible disruptive technology.

The annual tpmC coefficient of technical progress results are summarized in TABLE II. This would serve as a tool to forecast potential future performance expectations.

TABLE II
COEFFICIENT OF TECHNICAL PROGRESS 1/15/01

	Value
Average Coefficient of Technical Progress:	2.6403
Standard Deviation:	5.8409
Number of Observations:	59
95% Confidence Interval:	+/- 1.503

The statements of the trade journals concerning the arrival of Microsoft to the forefront of TPC are further enhanced by the fact that those DMUs with the greatest rates of change tended to be the larger open systems on proprietary hardware, see TABLE III [3, 5].

TABLE III
HIGHEST COEFFICIENTS OF TECHNICAL CHANGE

Date	Obs.	Hardware	Software	β_{annual}
9/15/2000	131	IBM eServer iseries 840-2420	AIX 4.3.3	46.44665
1/31/2000	165	Sun 6500 Cluster	Solaris 2.6	4.863074
3/1/1999	10	Bull Escalla EPC 2400 c/s	AIX 4.3.3	3.537029
3/1/1999	136	IBM RS6000 S80	AIX 4.3.3	3.291695

Those DMUs indicating the largest rate of efficiency frontier advancement are primarily proprietary hardware solutions of larger vendors. They correspond to those systems that were greatly affected by the release of the Win 2K submissions in 11/30/2000, 1/15/2001, and 9/30/2000, indicating that Win 2K and SQL Server may be a technological leap forward. By examining the migrating efficiency frontiers for a set of three dates, we observe that major changes in the higher end of the benchmarks were made with the advent of Windows 2000.

The lower elements of the rate of change, shown in Table IV indicate that the more mature models of the Microsoft NT and Windows 2000 seem to be holding their line fairly well. This could indicate quite a few things.

TABLE IV
LOWEST COEFFICIENTS OF TECHNICAL CHANGE

Date	Obs.	Hardware	Software	β_{annual}
3/6/00	114	HP NetServer LH 6000	MS NT 4.0	1.085119
9/30/00	28	Compaq ProLiant 8500-X700-96P	MS NT 4.0	1.111365
12/31/99	83	Dell PowerEdge 8450	Win2000	1.214633
5/6/99	48	Compaq ProLiant 5500-6/500	MS NT 4.0	1.274801

It must be noted that benchmarks are accurate at measuring a system's reaction to a benchmark environment and their efforts to mimic the real world often fall short. The TPC-C also has this downfall, and is essentially an effective measurement of how OLTP systems run TPC-C benchmarks. Therefore, this progress of technology relates to the ability of vendors to perform well on benchmarks. Much of the rapid increase may be associated to a better understanding of how to maximize the benchmarks for their own benefit, as well as the technical progress of the product. Much of the benchmarking has been improved through the loading of data into memory during the first parts of the tests resulting in faster scores for larger memory machines. By merely increasing the amount of memory able to be referenced by an operating system or application it is possible to increase one's scores.

4) Results with a Limited Time Horizon

As observed before, DMU 131 (6/1/2000) has an extremely high β when compared to the rest of the dataset. Such a high β could indicate that it is being compared to products representing a large technological advance when compared to the normal rate of progress. In this instance it is greatly affected by the two Microsoft entries to be made available on 11/30/2000 and 9/30/2000.

Several methods of dealing with this outlier are available. The first, and simplest procedure would be to simply delete DMU 131 from the calculations of Table II. For the sake of this exploratory study, we will demonstrate another alternative of examining technology progress relative to a slightly earlier time period before the two Microsoft entries. In this case, we are considering only products to be available

on 9/15/2000 or earlier. These results are shown in Table V, which demonstrates a more stable rate of change prior to the release of the larger disrupting Microsoft Windows 2000 configurations.

	Value
Average Coefficient of Technical Change:	1.7177
Standard Deviation:	0.2968
Number of Observations:	54
95% Confidence Interval:	0.07917




Figure 6 - Coefficients of Technical Change as of 09/15/2000

III. LIMITATIONS

A. Accuracy of the Data

Although TPC has been chartered to provide accurate and realistic representation of database performance, there does lie opportunity for “benchmark” scoring. That is tuning a database to increase performance in benchmark situations. The requirements for full-disclosure of results and the auditors of TPC provide insurance that the data is accurate. As time has progressed increased knowledge has allowed for a more optimal execution of the benchmark procedures. This is shown in part by an overall increase in top TPC-C results of nearly 10,000 times between 1992 and 1998, which can in part be attributed to real world performance increases and resolution of bugs discovered due to the benchmarks, but may well reduce the overall objectivity of the benchmarks [20].

B. Assumption of Technical Progress

One important thing to note is that this methodology assumes that technical progress is a linear regular progression over time. This approach would be a good fit for product categories that undergo incremental but significant innovation. It has been proposed that this may not be an accurate assumption as there are occasions where progress is not made as time progresses [18]. The concept of stasis lends itself well to the current model and as such may provide further future work.

IV. FUTURE WORK

A. Inefficient DMU Projection

The current model drops those DMUs that are inefficient upon release, since technical progress should be measured against those DMUs, which advance the efficiency frontier. An alternate method that may be more useful when considering a smaller number of DMUs is to project those inefficient DMUs to the efficiency frontier for future iterations. This would result in providing a larger number of DMUs in regions of heavier concentration and therefore market interest. This could enhance the reliability of the results obtained.

B. Alternate DEA Models

Other DEA models could be included such as Free Disposal Hull (FDH) to further determine the rate of change for the model. The FDH model of DEA simply says that the λ variables in the linear programming formulation of DEA need to be binary [22]. This eliminates the convexity assumption of DEA that implies that system configurations between actually systems should be possible. This simple model provides a good two-dimensional illustration as to how the methodology may be applied.

C. Future Technologies

Additional examination may also be given to the effect of newer technologies on the overall trends within the efficiency frontier. As was touched upon in the illustrations, outlier change tends to often be dramatic while the majority of the frontier tends to advance more slowly. Examining this trend may further assist in differentiating breakthroughs from regular incremental change. In addition, it may also be feasible to ascertain when a breakthrough is likely.

In regard to technology forecasting, much can be done to examine the effects of risk and varying inputs and outputs to ascertain the future reasonable performance characteristics of technologies.

D. Acceleration and Deceleration of Change

Of particular note may be the issue that technological progression is not necessarily a linear function and may go through waves (as experienced with the release of Microsoft Windows 2000 in this example). Overall, the goal should be to provide decision makers with reasonable pointers to indicate the feasibility of their future decisions. Other points of focus include identification of factors that may need to be increased or are lacking to push the technology frontier forward.

One thing assumed in this study is the constant rate of technological progress. Although convenient, it is a commonly accepted notion that technology does not advance at constant rates and is indeed prone to acceleration and deceleration over time, often corresponding to an S curve or similar pattern [12, 14, 21]. Future work included separation

and estimation of the annual rates of change on a yearly basis. This could then lead to a greater understanding of the causes of varying rates as well as being used as an early warning indicator for a maturing or stagnating product market. By providing these measures, decision makers may be able to identify potential for further investment.

E. Market Segmentation

While estimates of technical change may vary by segment, they are aggregated into a collective estimate and the current model does future efficiency frontier estimates to vary by segment of the efficiency frontier. Future work may extend the model to allow for quantitative analysis changes by market segment.

F. Validation

As the current data set is regularly updated by the TPC, further studies of more recent data sets may help further validate the premise that the model will allow for technology forecasting.

V. CONCLUSION

This method appears to provide some insight into the mapping of technical change in complex environments. This very simple example provides a straightforward means of examining the actual progression of technology over time and provides the users with the possibility of forecasting future performance of that technology.

REFERENCES

- [1] A. J. Alexander and J. R. Nelson, "Measuring technological change: aircraft turbine engines," *Technological Forecasting and Social Change*, vol. 5, pp. 189-203, 1973.
- [2] T. R. Anderson, S. Grosskopf, R. Färe, and X. Song, "Examining Moore's Law using data envelopment analysis," Submitted September 21, 2000 2001.
- [3] M. Apicella, "Uncovering database benchmarks," in *Infoworld*, vol. 22, 2000, pp. 63-64.
- [4] R. D. Banker, A. Charnes, and W. W. Cooper, "Some models for estimating technical and scale inefficiencies in data envelopment analysis," *Management Science*, vol. 30, pp. 1078-1092, 1984.
- [5] M. Biggs, "Understanding Benchmarks," in *Infoworld*, vol. 22, 2000, pp. 64.
- [6] A. Charnes, W. W. Cooper, and E. Rhodes, "Measuring the efficiency of decision making units," *European Journal of Operational Research*, vol. 2, pp. 429-44, 1978.
- [7] W. W. Cooper, L. M. Seiford, and K. Tone, *Data Envelopment Analysis: A Comprehensive Text with Models, Applications, References and DEA-Solver Software*. Dordrecht: Kluwer, 2000.
- [8] E. N. Dodson, "A general approach to the measurement of state of the art and technical advance," *Technology Forecasting*, vol. 1, pp. 391-408, 1970.
- [9] E. N. Dodson, "Measurement of the state of the art and technological advance," *Technological Forecasting and Social Change*, vol. 27, pp. 129-146, 1985.
- [10] R. Färe and S. Grosskopf, *Intertemporal Production Frontiers: With Dynamic DEA*: Kluwer, 1996.
- [11] R. Färe, S. Grosskopf, and C. A. K. Lovell, *Production Frontiers*: Cambridge, 1994.
- [12] S. Makridakis, S. C. Wheelwright, and V. E. McGee, *Forecasting Methods and Applications*. New York: John Wiley & Sons, 1983.
- [13] D. K. Peterson, P. E. Miller, W. A. Fischer, and R. W. Zmud, "Technology measurement and the appraisal of information technology," *Technological Forecasting and Social Change*, vol. 42, pp. 251-259, 1992.
- [14] A. L. Porter, A. T. Roper, T. W. Mason, F. A. Rossini, and J. Banks, *Forecasting and Management of Technology*. New York: John Wiley & Sons, 1991.
- [15] H. R. Rao and B. D. Lynch, "Hedonic price analysis of workstation attributes," *Communications of the ACM*, vol. 36, pp. 95-102, 1993.
- [16] D. Sahal, "The generalized distance measures of technology," *Technological Forecasting and Social Change*, vol. 9, pp. 289-300, 1976.
- [17] D. Sahal, "On the conception and measurement of trade-off in engineering systems: a case study of the aircraft design process," *Technological Forecasting and Social Change*, vol. 8, pp. 371-384, 1976.
- [18] D. Sahal, "Foundations of technometrics," *Technological Forecasting and Social Change*, vol. 27, pp. 1-37, 1985.
- [19] L. M. Seiford, "A bibliography for Data Envelopment Analysis (1978-1996)," *Annals of Operations Research*, vol. 73, pp. 393-438, 1997.
- [20] K. Shanley, "TPC Benchmarks," Transaction Processing Performance Council, http 1998 <http://www.tpc.org/information/benchmarks.asp>.
- [21] W. G. Sullivan and W. W. Claycombe, *Fundamentals of Forecasting*, 1st ed. Reston, Virginia: Prentice-Hall, 1977.
- [22] H. Tulkens, "On FDH efficiency analysis: some methodological issues and applications to retail banking, courts, and urban transit," *Journal of Productivity Analysis*, vol. 4, pp. 183-210, 1993.