Using stochastic cost frontiers to examine hospital efficiency

Andrew Street and Rowena van der Merwe
Centre for Health Economics, University of York

Paper presented at the Health Economics Study Group
Newcastle upon Tyne
5-7 January 2000

Work in progress; please do not cite without the authors' permission.

Table of Contents

Acknowledgements

The authors would like to thank Antonio Giuffrida, Peter Smith and Neil S derlund for their contributions to this work. All omissions and errors are the responsibility of the authors.

Abstract

In 1999 the Department of Health (DoH) developed a set of indices to compare the unit costs of NHS acute Trusts. The indices are based on an econometric model estimated using ordinary least squares (OLS) in which unit costs are explained by a range of factors, including casemix, specialty mix and hospital size. After making these adjustments, the remaining unexplained cost differences are deemed to represent inefficiency. This approach has been criticised for failing to account for random influences on production costs that are beyond the Trust's control. Using maximum likelihood (ML) techniques the residuals can be partitioned, with inefficiency assumed to be non-negative and the random (stochastic) element normally distributed around zero. In this paper the DoH model is re-specified as a stochastic cost frontier and re-estimated with alternative distributions of the inefficiency term. The efficiency rankings are well preserved across the stochastic cost frontier methods, but are less stable between the OLS and ML techniques.

1. Introduction

In its pursuit of greater efficiency and accountability of the acute healthcare sector, the Department of Health (DoH) has developed an econometric approach to measuring the unit costs of healthcare production (Söderlund and van der Merwe 1999). This approach is based on an econometric model in which unit costs are explained by a range of factors, including casemix, specialty mix and hospital size. After making these adjustments, the remaining unexplained cost differences are deemed to represent inefficiency. A Trust with a residual of zero is interpreted as displaying average technical efficiency, while those with negative (positive) residuals are deemed of above (below) average efficiency.

This approach has been criticised for failing to recognise that the unexplained component of the model is comprised of two elements: inefficiency and random influences on production costs that are beyond the Trust's control (Aigner et al. 1977). In econometric models this can be characterised by partitioning the residuals so that inefficiency is assumed to be non-negative and the random (stochastic) element is normally distributed around zero.

In this paper the economic and statistical logic behind this formulation is outlined (section 2) and the impact of estimating the DoH model using maximum likelihood (ML) rather than ordinary least squares (OLS) is explored with respect to the conclusions that may be drawn about relative hospital efficiency (section 3). Discussion of the results and concluding comments are offered in section 4.

2. Stochastic cost frontier analysis

The DoH has employed OLS methods to generate cost indices designed to assess the unit costs of acute Trusts, after taking account of various factors that explain differences in costs among Trusts. Full details of the approach are provided elsewhere (Söderlund and van der Merwe 1999).

In general terms, the econometric model is defined as:

$$y_i = \alpha + \boldsymbol{x}_i\boldsymbol{b} + \boldsymbol{e}_i \qquad i=1,\ldots,N$$

where y_i is the (total or unit) cost of production of the ith firm in either linear or logarithmic form; α is a constant; $\boldsymbol{x}_i$ is a k×1 vector of (transformations of the) input prices and output of the ith firm; $\boldsymbol{b}$ a vector of unknown parameters, and $\boldsymbol{e}_i$ is the error term. Under this formulation the error term is interpreted as representing inefficiency. In other words, the difference (residual) between observed unit costs and those predicted by the model is said to be due to inefficient behaviour.

This interpretation has been criticised for failing to recognise that at least some of the difference between observed and predicted costs may be random 'noise' arising from measurement error or unobservable heterogeneity. Stochastic cost frontier (SCF) analysis has been developed as a means to explore this possibility. The technique generates a stochastic error term and an inefficiency term by using the residuals from the estimated cost function. The approach decomposes the residual term into two parts:

$$\boldsymbol{e}_i = v_i + u_i \qquad\qquad i=1,\ldots,N$$

where v_i are the random (stochastic) elements, assumed independent and identically distributed (iid) with zero mean and variance σ_v^2, hence $v_i \sim N(0,\sigma_v^2)$.[a] u is a non-negative error term accounting for the cost of inefficiency in production. v_i and u_i are assumed to be independent (i.e. they have zero covariance).

The economic logic behind this specification is that the cost function is subject to two economically distinguishable random disturbances with different characteristics. v_i can be interpreted as representing stochastic events, such as unexpected winter pressure on beds arising from cold weather or a temporary local outbreak of disease. The u_i measures the degree of inefficiency in the ith hospital, usually represented as distance from the maximum possible cost given by the cost frontier. If allocative efficiency is assumed, the u_i are closely related to the cost of technical inefficiency. Technical inefficiency may arise from managerial slack (X-inefficiency), outmoded equipment, or inadequate staffing. If the assumption of allocative efficiency is not made, the u_i in a cost function incorporates both technical and allocative inefficiencies[b].

---

[a] This model collapses to an OLS model again when $\sigma_v^2 = 0$.

[b] Any failure in optimisation, whether technical or allocative, will show up as higher cost. The computation is dependent on the inputs chosen and whether they are allocatively efficient. Thus, a producer may be operating technically efficiently by a production function, but show up as inefficient with respect to a cost function. Therefore it has been argued that the interpretation of the one-sided error on the cost side as a measure of technical inefficiency is only appropriate if the measure is

The stochastic cost function is written as (Coelli 1996):

$$y_i = \alpha + \boldsymbol{x}_i\boldsymbol{b} + (v_i + u_i) \qquad i=1,...,N$$

In estimating the stochastic frontier, it is necessary to specify the distributional characteristics of the two components of the residual. Commonly it is assumed that v_i is normally distributed with a mean of zero, and u_i is non-negative. Under this specification the u_i defines how far the *i*th firm operates above the cost frontier. The inefficiency component u_i must be observed indirectly since the residual (y_i - α - $\boldsymbol{x}_i\boldsymbol{b}$) estimates $\boldsymbol{e}_i$ not u_i. The entire $\boldsymbol{e}_i$ (= v_i + u_i) could easily be estimated for each observation but for a long time an unresolved problem was how to separate it into its two components, v_i and u_i. Estimation of stochastic production and cost functions were accomplished by several groups of authors (Aigner et al. 1977; Schmidt and Lovell 1979), but because they were unable to solve the decomposition problem, none were able to obtain estimates of technical efficiency for each firm in the sample.

A solution to this problem was first proposed by Jondrow et al. (1982) by considering the expected value of u_i, conditional on (v_i + u_i)[c]. They specified the functional form of the distribution of the one-sided inefficiency component and derived the conditional distribution ($u_i \mid v_i + u_i$). The resulting residuals are decomposed to estimate the technical efficiency for each observation. Jondrow *et al* (1982) proposed a model that specified inefficiency with a half-normal distribution as follows:

$$E\left[u_i \middle| \boldsymbol{e}_i\right] = \frac{\boldsymbol{s}\,\boldsymbol{l}}{(1+\boldsymbol{l}^2)}\left[\frac{\boldsymbol{f}\left(\boldsymbol{e}_i\boldsymbol{l}/\boldsymbol{s}\right)}{\Phi\left(-\boldsymbol{e}_i\boldsymbol{l}/\boldsymbol{s}\right)} - \frac{\boldsymbol{e}_i\,\boldsymbol{l}}{\boldsymbol{s}}\right]$$

where Φ_u and Φ_v are the variance of the stochastic and inefficiency terms. $\lambda=\Phi_u/\Phi_v$ and captures inefficiency. Where $\lambda=0$, every observation would lie of the frontier (Greene 1993). $\boldsymbol{f}$ (.) and M (.) are, respectively, the probability density and cumulative distribution function of the standard normal distribution.

---

defined in terms of costs, rather than output. Thus one should measure efficiency by costs rather than outputs (Greene 1993). Consequently, inefficiency is often interpreted as “cost inefficiency”, the total of both technical and allocative inefficiency.

[c] Their solution was for the production function formulation of a stochastic frontier model with an error term (v_i - u_i).

Truncated normal and exponential distributions have also been proposed. For the truncated normal model, the explicit form is obtained by replacing the $\varepsilon_i\lambda/\sigma$ in the half-normal model with:

$$u_i^* = \frac{\varepsilon_i \lambda}{\sigma} + \frac{\mu}{\sigma \lambda}$$

If μ is not significantly different from zero, the model collapses to the half-normal. The expression for the exponential model, with a density function of the form $f(u)=\exp(-u/\sigma_u)/\sigma_u$, is (Jondrow et al. 1982):

$$E[u_i|\varepsilon_i] = (\varepsilon_i - \theta\sigma_v^2) + \frac{\sigma_v \phi[(\varepsilon_i - \theta\sigma_v^2)/\sigma_v]}{\Phi[(\varepsilon_i - \theta\sigma_v^2)/\sigma_v]}$$

The formulations for the various error distributions are as follows (Greene 1995):

Normal - half-normal model: $E[u] = \sqrt{2/\pi}\,\sigma_u$

$Var[u] = (\pi/2 - 1)\sigma_u^2$

Normal - truncated model: $E[u] = \mu + \sigma_u \lambda_u$

$Var[u] = \sigma_u^2 \left[1 - \lambda_u \left(\mu/\sigma_u + \lambda_u\right)\right]$

Normal - exponential model: $E[u] = 1/\theta$

$Var[u] = 1/\theta^2$

These formulations produce an unbiased but inconsistent estimator of u_i because, regardless of the sample size, the variance of the estimate remains non-zero (Greene 1993). The inconsistency of the estimator u_i is unfortunate in view of the fact that the purpose of the estimation is to approximate inefficiency. However, no improvements on this measure have yet been forthcoming in the literature for single-equation cross-sectional studies. The problem can be avoided if longitudinal data are available (Greene 1993).

This paper uses two econometric estimation packages, namely LIMDEP (Greene 1995) for the exponential distribution and FRONTIER (Coelli 1996) for the half-normal and truncated error distributions. Predictions of individual firm efficiencies are computed automatically in FRONTIER. The measure of efficiency relative to the cost frontier is defined as: $\text{EFF}_i = E(y_i|u_i, \mathbf{x}_i)/ E(y_i|u_i=0, \mathbf{x}_i)$, where y_i is the cost of the ith firm. EFF_i will take a value between one and infinity and can be defined as: $(\boldsymbol{x}_i\boldsymbol{b} + u_i)$ / $(\boldsymbol{x}_i\boldsymbol{b})$. While LIMDEP assumes a logarithmic functional form, FRONTIER allows the user to specify whether the dependent variable is in original units or logs. As the DoH model is in linear form, it is necessary to transform the efficiency scores produced by LIMDEP for the exponential distribution to make them comparable to those produced by FRONTIER.

3. Comparison of efficiency rankings

3.1. Data and model

The current paper reproduces analysis undertaken for the DoH, the results of which were circulated to acute hospital Trusts in early 1999. The first stage of the analysis involved calculating an index of actual to expected costs. This index was then regressed against series of independent variables that sought to explain variations in index scores. Data were compiled by the DoH from a variety of routine sources, including the Hospital Episodes Statistics, Trust Financial Returns, Trust Annual Accounts and the Hospitals Yearbook. The data for 1995/96 were made available on the NHSWeb [http://tap.ccta.gov.uk/doh/trustben.nsf].

The dependent variable, the casemix cost index (CCI), for hospital i is an index of actual over expected casemix weighted costs:

$$CCI_i = \frac{C_I}{[(IC*IP_i*H_i) / IP*H] + [\Sigma_j OP_{ij}*(OC_j/OP_j)] + [AE_i*AC/AE]} \quad (1)$$

where C_i is the cost of inpatient, outpatient and A&E care in hospital i; IC is the total cost of inpatient spells for all acute hospitals; IP_i is the number of inpatient (including day case) spells in hospital i; H_i is the Healthcare Resource Group (HRG) casemix index for hospital i; IP is the total number of inpatient spells for all acute hospitals; H is the average casemix index for all acute hospitals; OP_{ij} is the number of first outpatient attendances across all specialties in hospital i; OC_j is the total cost of outpatient attendances for all acute hospitals in specialty j; OP_j is the number of first outpatient attendances for all acute hospitals in specialty j; AE_i is the number of first

A&E attendances in hospital *i*; AC is the total cost of A&E attendances in all acute hospitals; and AE is the number of first A&E attendances in all acute hospitals.

The CCI is then regressed against a series of explanatory variables hypothesised to explain cost differences among hospital Trusts. Many of these adjustments attempt to account for the possibility that, even within an HRG, some hospitals will treat more costly patients. Hospital transfers, multi-episode spells, and the proportion of elderly or female patients are included to account for cost differences over and above the HRG casemix adjustment. In addition, allowance is made for possible cross-subsidisation between patient care and teaching or research which may not be adequately dealt with in the funding allocations, and for differences in local factor costs, assessed using the Market Forces Factor. The resulting index was termed the Casemix Costliness Cost Index (2CCI).

The variables are listed in Table 1 with basic descriptive statistics, for 218 hospital Trusts. These data exclude outliers.[d]

TABLE 1: DESCRIPTIVE STATISTICS

Variable key	Interpretation	Mean	Std Dev	Minimum	Maximum
Dependent variable					
CCI	Cost index	0.992	0.146	0.715	1.724
Independent variables					
TRANSIPP	Transfers in to hospital per spell	0.017	0.033	0.000	0.241
TRANSOPP	Transfers out of hospital per spell	0.021	0.017	0.000	0.125
EMERGPP	Emergency admissions per spell	0.348	0.091	0.020	0.748
FCEINPP	Finished consultant episode inter-specialty transfers per spell	0.019	0.014	0.000	0.114
OPNPP	Non-primary outpatient attendances per inpatient spell	2.909	0.855	0.000	7.847
EMERINDX	Standardized index of unexpected emergency admissions/total emergency admissions	0.059	0.020	0.016	0.258
EP_SPELL	Episodes per spell	1.068	0.117	0.785	1.661
HRGWTNHS	HRG weight, case mix index	94.212	20.918	72.018	242.028
PROP15U	Proportion of patients under 15 years of age	0.093	0.100	0.000	0.838
PROP60P	Proportion of patients 60 years or older	0.343	0.092	0.000	0.951
PROPFEM	Proportion of female patients	0.571	0.056	0.308	0.897
STUDENPP	Student whole time teaching equivalents per inpatient spell	0.001	0.001	0.000	0.012
RESEARPC	Percentage of total revenue spent on research (estimated 1995)	1.754	6.076	0.000	73.065
MFF_COMB	Market forces factor – weighted average of staff, land, buildings and London weighting factors	87.448	9.889	75.817	132.789

[d] Outliers were identified using the DFITS procedure, and applying a cut-off point where DFITS>3*$(k/n)^{0.5}$ where *k* is the number of parameters estimated and *n* the number of observations [Söderlund and van der Merwe 1999].

3.2. Results

Table 2 presents the original ordinary least squares (OLS) regression results for the 2CCI, as estimated by the DoH. These results are accompanied by the three stochastic frontier regression estimations, corresponding to the half-normal, truncated normal and exponential error distributions. The coefficients and the significance of the independent variables are similar across all specifications. This is to be expected, since both the OLS estimates (which provide the starting values for the iterations) and the maximum likelihood estimates used in the stochastic frontier regressions are consistent estimators (Greene 1993). The distribution parameters of both the half-normal and exponential models (***8*** and ***2*** respectively) are significant, suggesting that these models are an improvement to OLS estimation. In contrast, the truncated model yields a value for ***:*** that is not significantly different from zero and, as such, the model is equivalent to the half-normal.

TABLE 2: RESULTS FOR ALTERNATIVE SPECIFICATIONS OF THE 2CCI, 1995/96 DATA

Dependent Variable: CCI

Model	*Ordinary Least Squares regression*			*Stochastic frontier regression - half normal error distribution*			*Stochastic frontier regression – truncated normal error distribution*			*Stochastic frontier regression – exponential error distribution*		
	Coeff	**Std. Err.**	**t-ratio**	**Coeff**	**Std. Err.**	**t-ratio**	**Coeff**	**Std. Err.**	**t-ratio**	**Coeff**	**Std. Err.**	**t-ratio**
INTERCEP	0.38	0.22	1.73	0.40	0.20	1.96	0.37	0.20	1.89	0.36	0.20	1.85
TRANSIPP	1.38	0.41	3.37	1.51	0.33	4.65	1.51	0.32	4.67	1.53	0.43	3.53
TRANSOPP	-2.86	0.64	-4.44	-2.62	0.57	-4.60	-2.57	0.56	-4.61	-2.61	0.61	-4.25
EMERGPP	0.14	0.13	1.15	0.13	0.11	1.18	0.19	0.11	1.72	0.22	0.13	1.67
FCEINPP	0.88	0.59	1.51	1.29	0.52	2.46	1.26	0.50	2.52	1.25	0.70	1.78
OPNPP	0.06	0.01	5.33	0.05	0.01	4.94	0.04	0.01	4.57	0.04	0.01	4.01
EMERINDX	0.09	0.52	0.17	-0.28	0.40	-0.71	-0.39	0.40	-0.97	-0.48	0.62	-0.78
EP_SPELL	0.18	0.07	2.48	0.11	0.07	1.58	0.08	0.07	1.27	0.08	0.07	1.13
HRGWTNHS	0.00	0.00	2.67	0.00	0.00	2.25	0.00	0.00	2.54	0.00	0.00	2.47
PROP15U	-0.13	0.11	-1.14	-0.14	0.09	-1.58	-0.12	0.09	-1.32	-0.10	0.08	-1.30
PROP60P	-0.27	0.13	-2.08	-0.17	0.11	-1.52	-0.12	0.11	-1.04	-0.10	0.13	-0.76
PROPFEM	-0.16	0.20	-0.80	-0.10	0.18	-0.58	-0.02	0.18	-0.09	0.03	0.18	0.15
STUDENPP	12.99	7.79	1.67	14.12	7.46	1.89	15.54	4.39	3.54	16.64	9.82	1.69
RESEARPC	0.00	0.00	-0.87	0.00	0.00	0.16	0.00	0.00	0.65	0.00	0.00	0.46
MFF_COMB	0.00	0.00	2.73	0.00	0.00	2.59	0.00	0.00	2.36	0.00	0.00	2.06
μ							-0.53	0.58	-0.91			
θ										11.35	1.82	6.25
λ				2.90	0.47	6.11	2.03	0.86	2.36			
σ_v^2				0.0027			0.0044			0.0040		
σ_u^2				0.0230			0.0182			0.0078		
Log likelihood				189.7201			190.4361			190.5363		
R^2	0.5486			0.4827			0.4827			0.4827		
Adjusted R^2	0.5194			0.4470			0.4470			0.4470		
E[u]				0.1210			-0.2570			0.0881		
Var[u]				0.0131			0.0885			0.0077		
Mean efficiency	0.7810			0.8783			0.9010			0.9181		

It is evident from Table 2 that the mean level of efficiency is similar across the three types of stochastic error distributions, all around 90 percent. This is considerably higher than the mean level of efficiency in the OLS model of around 78 percent.

A graphical representation of each of the efficiency distributions deriving from each of the four specifications is presented in Figure 1. These figures are scaled such that the estimated efficiency increases along the horizontal axis. Figure 1a, deriving from the OLS model, can be interpreted as suggesting that few hospitals are efficient, with inefficiency distributed reasonably normally among the sample. The alternative specifications deriving from the stochastic frontier models imply that most hospitals are (or are close to being) relatively efficient. This corresponds to the relatively higher mean efficiency in the stochastic models.

The more pronounced negative skew of the exponential distribution (Figure 1d) implies that inefficiency is less widespread among hospitals than is assumed under the alternative specifications and a larger component of the error term is made up of random noise (or uncontrollable events) as opposed to inefficiency.[e]

[e] There has been some debate about whether or not this is an inevitable consequence of applying these distributional assumptions (Berger 1993; Cummins and Zi 1997).

FIGURE 1

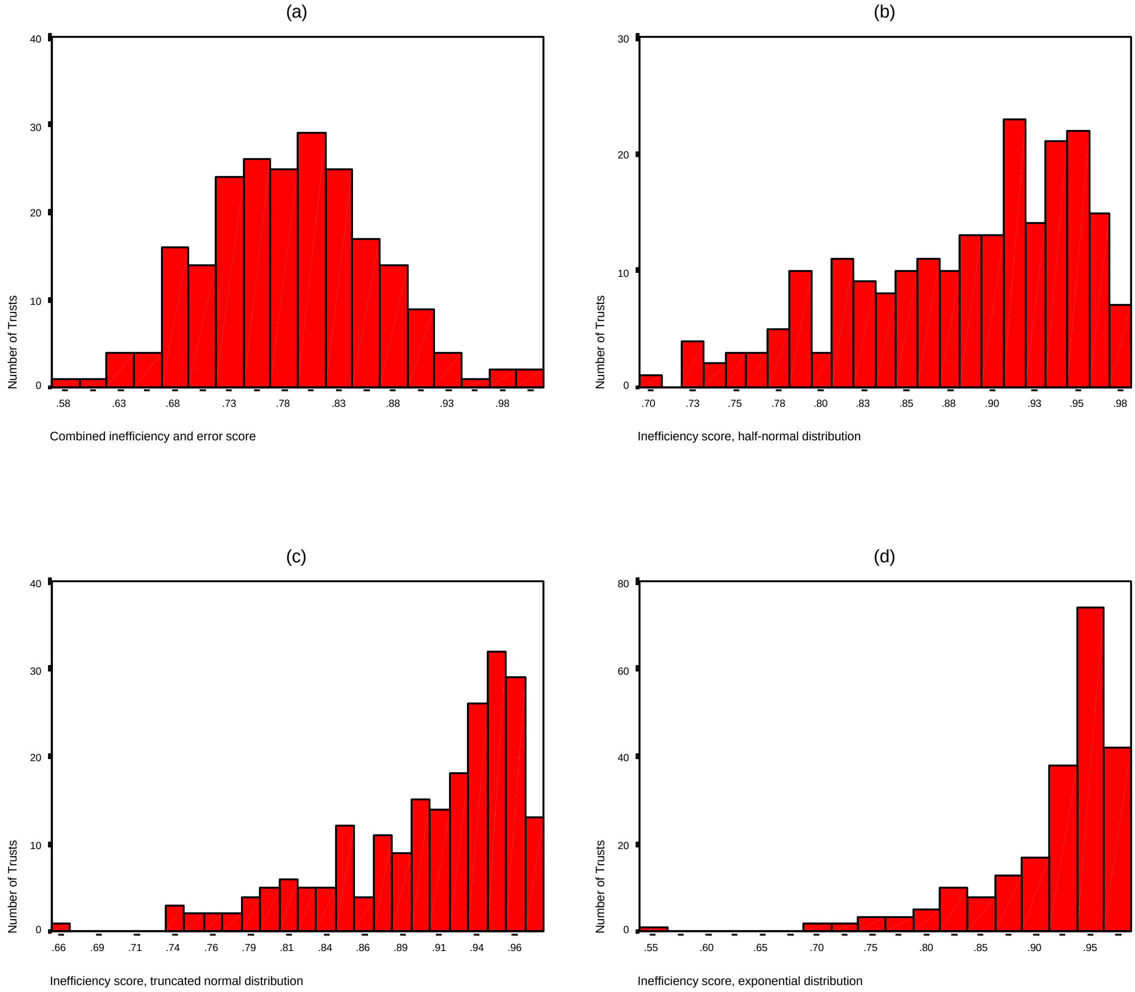

: HISTOGRAMS SHOWING (A) THE RESIDUAL FROM THE OLS MODEL; (B) THE INEFFICIENCY TERM FROM THE HALF-NORMAL MODEL; (C) THE INEFFICIENCY TERM FROM THE TRUNCATED MODEL; (D) THE INEFFICIENCY TERM FROM THE EXPONENTIAL MODEL (Note scales differ across graphs)

Table 3 shows the results for the correlation matrix of the efficiency scores and ranks from the OLS cost index (2CCI) and the stochastic cost frontiers estimations.

TABLE 3: RANK CORRELATIONS OF 2CCI COST INDEX AND STOCHASTIC COST FRONTIER EFFICIENCY SCORES AND RANKS

	2CCI	2CCI-half	2CCI-trunc	2CCI-exp
SCORES				
2CCI	1.000			
2CCI-half	0.883	1.000		
2CCI-trunc	0.848	0.983	1.000	
2CCI-exp	0.791	0.918	0.971	1.000
RANKS				
2CCI	1.000			
2CCI-half	0.908	1.000		
2CCI-trunc	0.897	0.996	1.000	
2CCI-exp	0.910	0.979	0.989	1.000

Note: 2CCI-half, 2CCI-trunc and 2CCI-exp are half normal, truncated normal and exponential error distributions, respectively, of stochastic frontier regressions of 2CCI variables.

The raw scores and ranked scores of the efficiency estimates derived from the stochastic models (2CCI-half, 2CCI-trunc and 2CCI-exp) are almost perfectly correlated. In other words, relative efficiency estimates remain stable regardless of the error distribution that is imposed. There is lower correlation between the efficiency scores derived from the OLS cost index and the stochastic frontier models. This implies that the alternative specifications of the inefficiency term under the stochastic models will have minimal effect on the relative ranking of the efficiency of hospitals, but that changing the specification from an OLS model to a stochastic model will have greater impact.

This point is illustrated in Figures 2 and 3. These figures plot the rank of hospitals on a scale from 1 to 218, with 1 indicating the most efficient hospital. The length of each line indicates the impact of changing the model specification on each hospital’s ranking. In figure 2, rankings derived from the three stochastic frontier models are compared for acute London hospitals (comparisons are similar within other family groups). The lines are fairly short, suggesting that the alternate distributions have limited impact on each hospital’s relative ranking. Figure 3 represents the change in ranks across the OLS and stochastic frontier models. Lines are generally slightly longer. This suggests that the choice of a OLS rather than stochastic model has greater influence on conclusions about relative efficiency than does the choice of which distribution should be adopted within the stochastic frontier framework.

FIGURE 2: CHANGE IN EFFICIENCY RANKING ACROSS STOCHASTIC COST FRONTIER FORMULATIONS FOR ACUTE LONDON TRUSTS

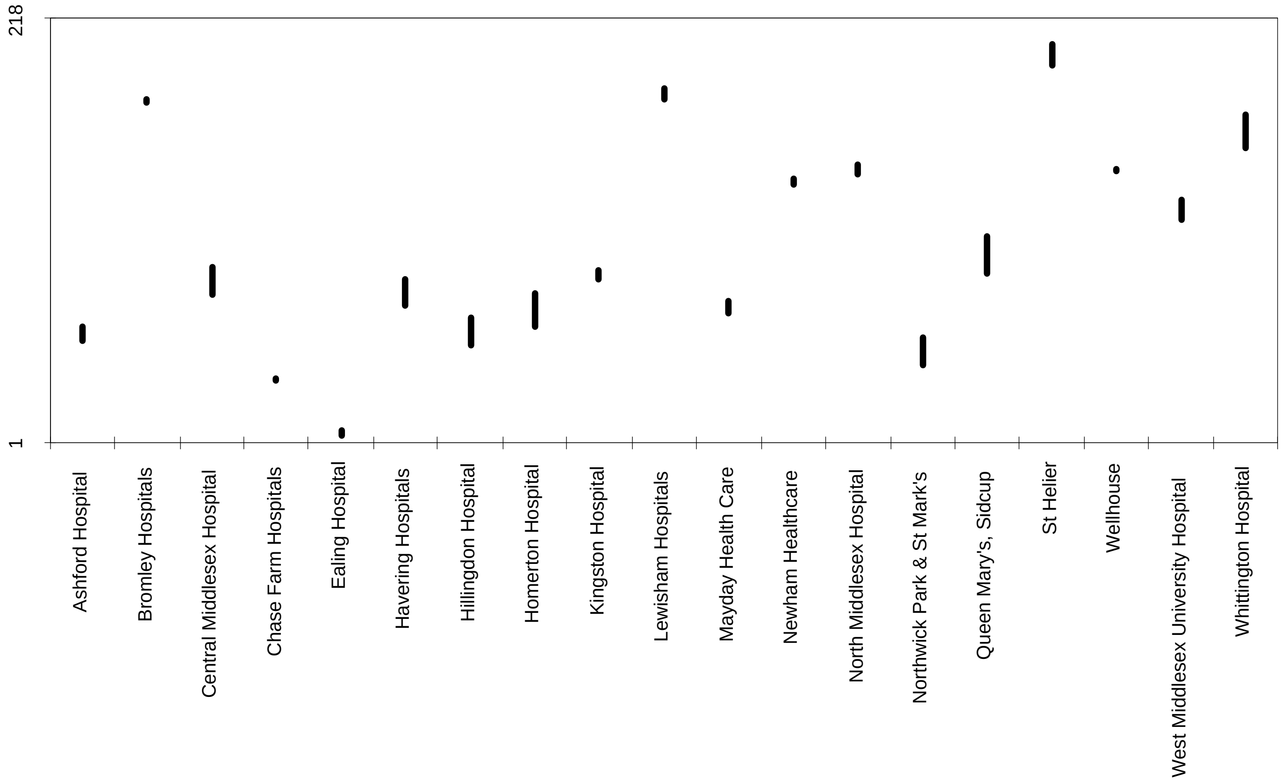

FIGURE 3: CHANGE IN EFFICIENCY RANKING USING OLS AND STOCHASTIC COST FRONTIER FORMULATIONS FOR ACUTE LONDON TRUSTS

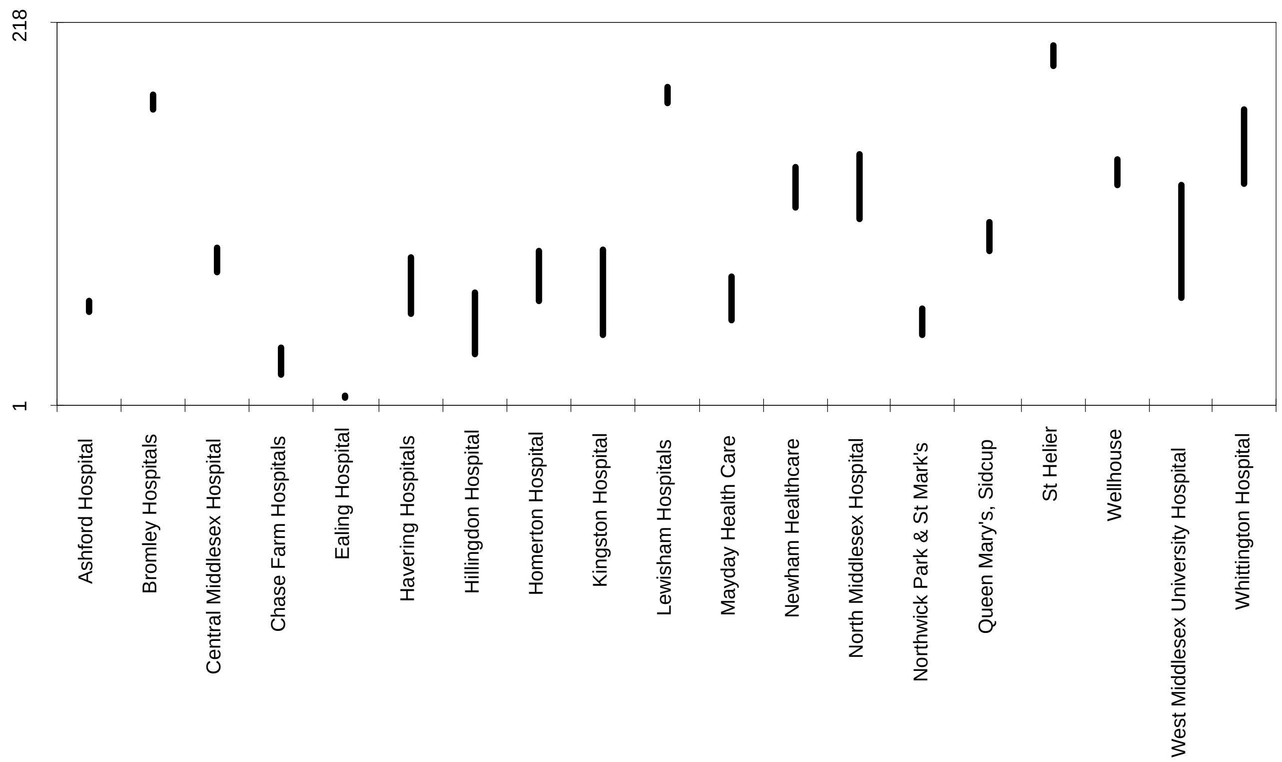

4. Discussion

This paper compares the unit costs of acute Trusts in the English NHS using alternative econometric techniques. The main findings suggest that specification as a stochastic cost frontier rather than an OLS model can lead to substantial differences in the estimated cost efficiency values, but that the efficiency rankings of Trusts are reasonably well preserved. Within the stochastic cost frontier framework, the choice of error distribution for the inefficiency term has a minimal impact on the estimate of mean efficiency and on the relative ranking of Trusts.

This consistency compares with the greater effect on the estimates of relative hospital performance of adopting a mathematical programming approach, such as Data Envelopment Analysis (DEA), rather than an econometric technique (van der Merwe forthcoming). These differences may stem from inadequate specification of the econometric model (which was taken as given here, as it was devised by the DoH) or because DEA fails to distinguish inefficiency from random noise. This paper has suggested that random shocks and measurement error may be an important component of residual variance.

The analysis has been constrained by the cross-sectional data. If panel data are available it would be possible to adopt a 'distribution free' method (Berger 1993; Schmidt and Sickles 1984) in which the inefficiency term is assumed constant over time. Construction of a panel for NHS Trusts requires consistency in data definitions and counting practices across years, or some means of dealing with inconsistencies. However, Trusts have continually revised their approaches to counting activity in order to realise annual productivity targets (Radical Health Statistics Group 1992; Radical Health Statistics Group 1995). Moreover, longitudinal estimation would need to make allowance for the fact that the unit of analysis has not remained constant over time, as a result of Trust mergers. Nevertheless, panel data estimation of NHS hospital costs may be a fertile subject of future research.

5. References

Aigner, D., Lovell, C.A.K. and Schmidt, P. 1977. 'Formulation and estimation of stochastic frontier production function models'. Journal of Econometrics 6: 21-37.

Berger, A.N. 1993. "Distribution-free' estimates of efficiency in the US banking industry and tests of the standard distributional assumptions'. Journal of Productivity Analysis 4: 261-92.

Coelli, T. 1996. 'A guide to FRONTIER version 4.1: a computer program for stochastic frontier production and cost function estimation' : Centre for Efficiency and Productivity Analysis, University of New England.

Cummins, J.D. and Zi, H. 1997. 'Measuring cost efficiency in the US life insurance industry: econometric and mathematical programming approaches' . Philadelphia: The Wharton School, University of Pennsylvania.

Greene, W.H. 1993. 'The econometric approach to efficiency analysis, in' in Fried, H.O., Lovell, C.A.K. and Schmidt, S.S. (eds.) The measurement of productive efficiency: Techniques and applications. New York: Oxford University Press.

Greene, W.H. 1995. Limdep Version 7.0 User's Manual. Castle Hill: Econometric Software, Inc.

Jondrow, J., Lovell, C.A.K., Materov, I.S. and Schmidt, P. 1982. 'On the estimation of technical inefficiency in the stochastic frontier production function model.'. Journal of Econometrics 19: 233-8.

Radical Health Statistics Group 1992. 'NHS Reforms: the first six months - proof of progress or a statistical smokescreen'. British Medical Journal 304: 705-9.

Radical Health Statistics Group 1995. 'NHS 'indicators of success': what do they tell us?'. British Medical Journal 310: 1045-50.

Schmidt, P. and Lovell, C.A.K. 1979. 'Estimating technical and alloctive inefficiency relative to stochastic production and cost frontiers.'. Journal of Econometrics 9: 343-66.

Schmidt, P. and Sickles, R.C. 1984. 'Production frontiers and panel data'. Journal of Business and Economic Studies 2: 299-326.

Söderlund, N. and van der Merwe, R. 1999. 'Hospital benchmarking analysis and the derivation of cost indices' . York: Centre for Health Economics, University of York.

van der Merwe, R. forthcoming. 'Alternative methods to examine hospital efficiency: Data Envelopment Analysis and stochastic frontier analysis' . York: University of York.