

MAINTENANCE

JONATHAN E. LEIGHTNER*

College of Business Administration, Augusta State University

Maintenance and production interact. The ideal way of accounting for this interaction, when estimating production functions, is by picking the temporal length of observations so that they embed integer multiples of the production–maintenance cycles for all inputs. In contrast to labor and land, the production–maintenance cycles of capital sometimes vary tremendously in temporal length, which can make it impossible to implement the ideal method of accounting for the interaction between maintenance and production. This paper empirically tests four second best methods of accounting for maintenance, when the ideal method is impossible. The output elasticities of all inputs (not just the input undergoing maintenance), which emerge from these tests, vary tremendously. This implies that the way that maintenance is incorporated into the analysis (including the standard approach of ignoring maintenance) drastically affects the profit maximizing combinations of inputs derived from production function estimations.

JEL classification codes: D24, D92, L20, L94

Key words: maintenance, efficiency, electricity, production and cost functions

I. Introduction

Maintenance is a phenomenon which permeates almost every type of

* I would like to thank two anonymous referees for their helpful comments, especially on focusing this paper. I also would like to thank Knox Lovell, Mike Dunn, and James Peach for their help with the data analysis and with earlier drafts of the paper. I appreciate those who helped acquire the data and who helped explain Mae Moh. These people include the Yonok Research Institute staff, especially Yupa Phusahas and Watcharapong Pienchob, Walaiporn Mohr, and countless people at EGAT and Mae Moh, especially Wittaya Kocharak, Ramej Pathumasut, Surachai Chaipurksayanon, Montri Suwanmontri, and Somsong Khakhay.

productive activity and is required for most inputs. Automobiles must be stopped to change their oil. Huge coal fired electricity generators must be shutdown and cooled, to better clean the ash out of them and to replace burned out parts. The farmer needs to stop planting and harvesting in order to weed, fertilize, and insecticide his land. Labor needs to stop in order to eat and to rest. Notice that in all these examples, maintenance results in normal production being stopped (or at least slowed) and inputs (which are currently not undergoing maintenance) being diverted from normal production and into the maintenance activity.

Maintenance's effects on outputs and inputs directly affect the estimated marginal products of inputs which, in turn, affect the calculated profit maximizing combinations of inputs. In spite of these effects, economists ignore maintenance when estimating production functions and cost curves. Extensive searches of the economics literature produces no results – I found no articles on maintenance's effects on the estimation of production functions and cost curves.¹

The following three sections of this paper correspond to my three primary objectives. Section 2 explains the ideal method of accounting for maintenance when estimating production functions and cost curves. Section 3 presents second best methods for accounting for maintenance when the ideal method is not practical. Section 4 shows empirically how maintenance affects the estimation of production functions. Throughout this paper, I will demonstrate

¹ I have found only one "borderline" exception to this hole in our literature. Rothwell and Rust (1995) use a dynamic programming (DP) model to predict shutdowns of nuclear power plants. However, they do not separate shutdowns due to refueling, maintenance, and/or false safety alarms (p. 19). Furthermore, they "found that most of the downtime is spent in periodic refueling outages" (p. 37). Although managers are likely to seize the opportunity presented by refueling to conduct preventative maintenance, the primary purpose of these shutdowns is refueling and the amount of maintenance conducted is unknown. Moreover, they deleted "the 56 major problem spells (defined as outages lasting more than 9 months) from our estimation sample" (p. 50). But these are the down times most likely to involve major maintenance. Rothwell and Rust deal more with refueling than maintenance.

my major points by using data from EGAT Mae Moh, the power plant that generates one third of Thailand's electricity consumption. Specifically, I will show why the ideal method of correcting for maintenance is impossible when analyzing EGAT Mae Moh's production. I will then present the empirical results of utilizing the four different second best approaches when estimating Cobb-Douglas and Translog production functions for EGAT Mae Moh. Section 5 is the conclusion.

The empirical results produce two important conclusions. First, when production function estimates that are not corrected for maintenance are compared to those corrected for maintenance, it is found that correcting for maintenance noticeably affects the estimated output elasticities for all inputs, not just the input undergoing maintenance. Second, in comparison to not correcting for maintenance, each input's elasticity was increased by some second best methods and decreased by other second best methods.

II. The Ideal Method to Correct for Maintenance

Since production and maintenance affect each other it is best to model these two activities together. Furthermore, the best way to model production-maintenance is by picking time periods of the optimal temporal length for the analysis. The importance of the temporal length of the observations can be seen by comparing the relationship between maintenance and production when (1) an extremely short time period is picked to when (2) an extremely long time period is picked. If the length of the time period is equal to (or less) than the time it takes to conduct the maintenance, then a negative relationship between output and maintenance will be found. This is because the input undergoing maintenance usually must be shutdown to conduct the maintenance and other inputs (especially labor) must be diverted from normal production in order to conduct the maintenance. In contrast, if the length of the time period is equal to the life span of the input, then a positive relationship will probably be found between maintenance and output. Clearly, when

maintenance is involved, the length of the time period used in the analysis is crucial.

The optimal length of the time period for the analysis will depend on the type of maintenance. Consider the following types of maintenance: (1) enhancing the marginal products of traditional inputs and (2) extending the life span of traditional inputs.² The man who turns off his band saw in order to change its dull blade does so in order to increase the subsequent productivity of both himself and the band saw. This act does not extend the life of the band saw. In contrast, if the man shuts off the band saw in order to oil it, the productivity of both himself and the band saw in the immediate future is probably not increased, but the longevity of the band saw is.

If the primary purpose for maintenance is to extend the life span of an input, then the optimal method of correcting for maintenance is by picking time periods which correspond to the entire life span (or whole integer multiples of the life span) of that input. Unfortunately, this is impractical for most studies because the number of observations would be excessively reduced. For example, in October 1994 Mae Moh power plant consisted of 11 generating units of three different technological vintages. Generating units 1-3 are 75 MW units built between 1978 and 1981. Generating units 4-7 are 150 MW units built between 1984 and 1985. Generating units 8-11 are 300 MW units built between 1989 and 1991. The average life span of the 75 MW units is 30 years. I have 117 monthly production data observations for Mae Moh between December 1984 and August 1994. Therefore, if maintenance is primarily longevity producing, my 117 observations would be reduced to one third of a single observation (or less) if I used the optimal time frame. Clearly the optimal method is impractical.

² It is possible to think of maintenance as depreciation forestalling. Depreciation also comes in these same two forms. One type reduces the immediate marginal products of the depreciating inputs and the other type shortens the life of the input. This distinction seems to be related to a strong debate in the depreciation literature. See Jergenson (1996) and Triplett (1996) for surveys of this literature and how it relates to this debate.

If, however, the primary purpose for maintenance is increasing the marginal productivity of the input immediately following the maintenance, then the optimal method for correcting for maintenance is by picking time periods for the analysis which correspond to complete production-maintenance cycles (or whole integer multiples of these cycles). Since production-maintenance cycles are shorter than the entire life span of an input, correcting for “marginal product enhancing” maintenance reduces the number of observations less than correcting for “longevity producing” maintenance.

Most of the maintenance conducted on labor is “marginal product enhancing.” Labor is given coffee breaks and lunch breaks, not so labor will live longer, but so that labor will be more productive after the breaks. When a laborer is given a medical leave of absence to have surgery to remove cancer, then perhaps longevity producing maintenance is being conducted. However, this type of maintenance usually is arranged by the laborer, not by his employer. I believe that the maintenance of labor would skew the results of production analysis if time periods shorter than one day were used. For example, if hourly data was used, then output would fall drastically during the lunch hour, rise immediately following breaks and fall just prior to breaks.³

Assuming, with good reason, that the primary purpose behind maintenance on land is marginal product enhancing, then the optimal time frame for the analysis of agricultural production is a year. If monthly data were used to analyze agricultural production instead of annual data, then the empirical results would be skewed. For example, large amounts of inputs would be used during planting months with no output recorded.

Just as the hourly analysis of the productivity of labor and the monthly analysis of the productivity of agricultural land is unreasonable, likewise the

³ It could be argued that complete production-maintenance cycles for labor should be based on weekly data in order to account for the rest which occurs during weekends. Perhaps, but this does not change the major point being made, which is that time periods for observations should be picked which embody integer multiples of the production-maintenance cycles of all inputs.

analysis of the productivity of capital which does not use complete production-maintenance cycles is also unreasonable. Ideally, each observation used in production analysis should include integer multiples of the complete production-maintenance cycles for all inputs involved— land, labor, and capital. It is relatively easy to get complete production-maintenance cycles (or integer multiples thereof) for labor and land because the production-maintenance cycles of these inputs are extremely consistent. Unfortunately, the production–maintenance cycles of capital are often not as temporally consistent as the cycles for land and labor.

For example, Mae Moh's generating unit 1 underwent 9 complete maintenance-production cycles between October 1979 and October 1994. Thus, if I used time periods which correspond to unit 1's complete production-maintenance cycles, then 180 monthly observations would be reduced to 9 production-maintenance cycle observations. Furthermore, each of these production-maintenance cycle observations would be of different temporal lengths. For generating unit 1, these cycles lasted an average of 573 days but the length of these cycles had a standard deviation of 181—almost 1/2 of a year.⁴ Unit 1's longest cycle lasted 948 days and its shortest cycle lasted 395 days. Therefore, Mae Moh's maintenance-production cycles are not relatively stable in length. Furthermore, the variations in the production-maintenance cycles of Mae Moh's eleven generators were not synchronized, making it impossible to find time periods of whole integer multiples of the production-maintenance cycles of all eleven generators. The ideal method of correcting for maintenance is impossible for Mae Moh.

Ceteris paribus, the empirical problems resulting from ignoring maintenance on capital increase as the number of machines increase, as variance in the temporal length of maintenance-production episodes increase, as the number of emergency maintenance episodes increase, and as the number

⁴ The average length of the maintenance part of these cycles was 46.7 days with a standard deviation of 9.53 days. The average length of the production part of these cycles was 526 days with a standard deviation of 177 days.

of technological vintages increase. Because most large firms use numerous machines of different technological vintages and because the temporal length of the maintenance-production cycles of each machine is not consistent, the ideal method of correcting for the maintenance of capital is impossible for most large firms. When researchers have estimated production functions for these firms, the maintenance of capital has been ignored. The effects of ignoring the maintenance of capital when estimating production functions is analogous to using hourly data when labor is involved or monthly data when agricultural land is involved.

III. Second Best Methods for Accounting for Maintenance

If the ideal way of modeling production and maintenance together is impractical, then modeling “production net of maintenance activities” may be the next best option. For Mae Moh the traditional output producing activities consist of combining the traditional inputs of capital, labor, and lignite in order to produce the traditional output, electricity. Separating Mae Moh's maintenance on its generators from these traditional output producing activities can be conducted in at least five ways: (Model 1) all observations contaminated by maintenance activity could be eliminated from the sample, (Model 2) the capital input can be adjusted down for maintenance, (Model 3) the capital and labor inputs can be adjusted down for maintenance, (Model 4) the electricity generated and the lignite consumed can be adjusted up to compensate for maintenance, and (Model 5) maintenance could be added as a second output.

“Production net of Maintenance Model 1” needs no further explanation. In Model 2, gross capital is replaced by net capital in a production function framework. The rationale behind this is that capital shut down for maintenance is not producing and should not be considered when estimating production functions. Net capital, for a given technological vintage “*i*”, was calculated by using equation 1:

$$\text{Net Capital } i = \sum_{j=k}^n [G_j/D] * [MW_i] \quad (1)$$

Where “G” is the number of days in the month that the unit was not shutdown for maintenance, “D” is the total number of days in that month, “MW_i” is the generating capacity of that technological vintage, and “j” is the unit number (“j” would range from 1 to 3 for the 75 MW units, 4 to 7 for the 150 MW units, and 8 to 11 for the 300 MW units).

In addition to using net capital instead of gross capital, “Production net of Maintenance Model 3” also replaces (gross) labor with net labor (total labor minus labor conducting maintenance). The rationale for this is if labor is being absorbed by the maintenance activity, then it also should not be included when estimating production functions. If labor used for normal production is separated from labor used for maintenance in the data, then Model 3 is easy to implement and it is preferred over Model 2. Unfortunately, my data on Mae Moh (like most data sets) does not separate normal production labor from maintenance labor, making it necessary to approximate “net labor.” “Net labor” is approximated by multiplying gross labor by the ratio of net capital to gross capital. Using this approximation for net labor assumes that labor can easily be used for both maintenance and production and that the percent of labor absorbed by maintenance is proportionate to the percent of capital shut down for maintenance. If these assumptions hold, then model 3 is better than model 2. If however, maintenance labor is specialized and relatively constant while overall labor fluctuates primarily with output or lignite used, then model 2 is better than model 3.

Production net of Maintenance Model 4 is accomplished by using gross capital and gross labor, but adjusting electricity and lignite consumption upward by multiplying them by the ratio of gross capital to net capital. The rationale for this is that electricity generation and lignite consumption would have been proportionally higher if maintenance had not shut down some of

the capacity. Production net of Maintenance Model 5, treating maintenance as a second output, can not easily be used in a production function framework because it would involve two dependent variables. However, it could be done in a cost function analysis. Unfortunately, I do not have complete cost data for Mae Moh.

IV. Empirical Results

I have maintenance, output, and input data on EGAT Mae Moh, the company in Thailand which generates 30% of Thailand's electricity consumption. The maintenance data for Mae Moh spans from October 1979 to October 1994. October 1994 is just prior to Mae Moh adding two more 300 MW units with desulfurization plants (for pollution abatement purposes) attached and retrofitting desulfurization plants to units 8-11. Therefore, October 1994 is just prior to another jump in capacity and to a change in technology. The maintenance data consists of which generating unit was shut down, the date it was shut down, and the date maintenance was finished.

I have monthly data on electricity generated, lignite consumed, and capital for October 1984 to August 1994 for the three 75 MW generating units combined (#1-3) and for the four 150 MW generating units combined (#4-7). The monthly data on the four 300 MW generating units (#8-11) begins in April 1989, when the first of these units came on line, and also extends to August 1994. All of this data is in physical units -gross generation of electricity is in MkWh, lignite is in thousands of tons, and capital is the generating capacity of the boiler, turbine, generator, and accessories. I have labor data consisting of the total number of permanent employees working in units 1-11 combined for each September of 1988 to 1994 and for each December of 1985 to 1994. To convert this labor data into monthly data, I assume a uniform monthly change in the data I do have. I allocate the resulting monthly labor

data to generating units 1-3, 4-7, and 8-11 in proportion to their share of total lignite consumed.⁵

⁵ A data inconsistency exists for generating unit 8 in July 1990. According to the maintenance data, unit 8 was shut down for maintenance from June 16, 1990 to August 15, 1990; however, the electricity data indicates that unit 8 generated 220 MWh of electricity during July of 1990 (a relatively large amount of generation for unit 8) when it supposedly was shut down. Therefore, I suspect that the wrong generating unit was given in the maintenance data (I know it was not one of the 300 MW units, but it could have been one of the 150 or 75 MW units). Therefore I eliminated this maintenance observation when conducting the analysis.

Getting the lignite into the generators, the ash out of the generators, and maintenance are the most labor-absorbing tasks performed by permanent labor in the Operations Division of Mae Moh. Maintenance may lead to a negative relationship between lignite and labor, but getting the lignite into the generators and the ash out of the generators should create a positive relationship between lignite and labor. I regressed permanent labor against lignite and the amount of capacity shut down for maintenance. After correcting for third degree autocorrelation, the Durbin-Watson statistic indicated that the hypothesis of autocorrelation can be rejected at a 95% confidence level. All coefficients were of the expected signs and the t-statistics indicate that the maintenance coefficient is not significantly different from zero at a 90% confidence level, but the lignite coefficient was. I interpret these results as supporting my allocation of labor according to lignite used.

Alternatively, I could have ignored the statistical insignificance of the maintenance coefficient and allocated labor to lignite and maintenance via the coefficients in this estimation. I chose not to do this because I interpret the third degree autocorrelation as indicating that variations in labor are not driven by monthly changes in lignite consumption or maintenance, but rather by longer run considerations and, in the long run, the amount of maintenance is (to a large extent) driven by the amount of lignite consumed since the last period of maintenance was performed. There was no good way to allocate the labor data to the various technology vintages, I have picked what I believe is the best of several bad options.

When conducting the analysis, I reintroduce the notion that maintenance may affect the amount of labor available for production in Production net of Maintenance Model 3. This model uses net labor and net capital when estimating production functions and frontiers with disaggregate data. To calculating net labor for model 2, I take aggregate labor and multiply it by aggregate net capital over aggregate gross capital in order to calculate aggregate net labor. I then allocated aggregate net labor to the different vintages according to lignite consumption.

I use 2 different models of production in which three inputs (lignite, capital, and labor) are used to produce one output (electricity): (1) Cobb-Douglas Production Function and (2) Translog Production Function. The Cobb-Douglas production function model is imposed when the following equation is estimated using ordinary least squares: where “ln” is the natural log, “y” is output, and “x” is used for inputs, and “n” is the number of inputs.

$$\ln y = \alpha_o + \sum_n \alpha_n (\ln x_n) \quad (2)$$

Under the Cobb-Douglas functional form, the marginal product elasticity (ϵ) for input “n” is α_n and returns to scale (RTS) are equal to $\sum_n \alpha_n$ for all observations. The Translog production function model is imposed when the following equation is estimated:

$$\ln y = \alpha_o + \sum_n \alpha_n (\ln x_n) + \sum_i \sum_{j, i \neq j} \alpha_{ij} (\ln x_i) (\ln x_j) + \sum_i \alpha_{ii} (1/2) (\ln x_i)^2 \quad (3)$$

Under the Translog functional form, ϵ and RTS vary from observation to observation. The translog ϵ for input “n” for a given observation equals

$$\epsilon_n = \alpha_n + \sum_j \alpha_{nj} (\ln x_j) \quad (4)$$

and RTS for that observation equals

$$\sum_i \alpha_i + \sum_i \sum_j \alpha_{ij} (\ln x_j) \quad (5)$$

In order to compare the “Production net of Maintenance Models” described in Section 3, I need samples in which gross capacity varies. This requirement leaves me with two choices: (1) using an aggregate sample of 117 observations spanning from December 1984 to August 1994 or (2) using a disaggregate sample of 65 observations for the 300 MW generating units (8-11) combined. The disaggregate sample spans from when the first 300 MW unit came on

line (April 1989) to the end of my data set (August 1994). Gross capacity does not vary for the 75 MW units (1-3) or for the 150 MW units (4-7) between December 1984 and August 1994; therefore disaggregate samples based on these vintages were not possible. The aggregate sample is created by summing the electricity generation and quantity of each input over all the technological vintages for each month in the sample.

The advantage of the aggregate sample is its relatively larger size; however, it has the major disadvantage of technological heterogeneity which gets progressively worse over time. A Dickey-Fuller test on the aggregate sample indicated that the appropriate way to detrend this data is via adding a trend term (in contrast to first differencing) when estimating a production function. In spite of adding a trend term, I had to correct for 8th degree autocorrelation using a Cochrane-Orcutt technique before the Durbin-Watson statistic indicated that there was only a 5% chance of continued autocorrelation in the production function estimates. When estimating the production function with the disaggregate data, the Dickey Fuller test indicated that first differencing was the preferred method of detrending the data and the resulting Durbin-Watson statistics indicated that all the disaggregate estimates needed, at most, a first degree correction for autocorrelation. I interpret the depth of the autocorrelation in the aggregate sample versus, at most, just first degree autocorrelation in the disaggregate sample as indicating that technological heterogeneity, which gets worse over time and which is inherent in the aggregate sample, is a serious problem.⁶ Thus, I only present the estimation results for the disaggregate data.

Table 1 presents the Cobb-Douglas and Table 2 the Translog production function empirical results. On both tables, Panel A presents the elasticities, Panel B, the t-statistics, and Panel C, the R-Bar Squares. The R-Bar Squares

⁶ The first degree autocorrelation in the disaggregate sample could be due to seasonal weather patterns trapping different amounts of SO₂ in the local Mae Moh valley, which then required seasonal adjustments in pollution abatement efforts (Leightner and Lovell, 1998).

for both the Cobb-Douglas and Translog are 0.99 indicating that these estimations explain 99 percent of the variation in output. Under the Cobb-Douglas functional form, the elasticities are the estimated coefficients; whereas, under the translog functional form, the elasticities are calculated using equation 4 above.

Table 1. Cobb-Douglas Production Function Empirical Results*

	0. No correction	1. Observations deleted with maintenance	2. Net capital	3. Net capital and net labor	4. Adjusted output & lignite
I. Elasticities					
ϵ lignite	0.950	0.907	0.939	0.973	0.975
ϵ capital	-0.038	-0.073	-0.043	-0.039	-0.031
ϵ labor	0.106	0.127	0.119	0.088	0.079
RTS	1.018	0.961	1.016	1.022	1.023
II. T-statistics					
ϵ lignite	23.288	12.725	21.839	27.184	32.262
ϵ capital	-1.430	-2.210	-1.595	-1.288	-1.208
ϵ labor	2.143	1.588	2.258	1.778	1.982
III. R Bar²					
	0.999	0.998	0.999	0.999	0.999

* ϵ = marginal product output elasticity and RTS = returns to scale. The data was first differenced because Dickey-Fuller tests indicated that 1st differencing was the correct way of detrending the data (instead of adding a trend term). The Durbin-Watson Statistics indicated that the production net of maintenance models 0 and 2-4 did not need to be corrected for Autocorrelation, but model 1 did. For model 1, a correction for 1st degree autocorrelation using a Cochrane-Orcutt technique, resulted in a Durbin-Watson statistic which indicated that no further correction for autocorrelation was needed. After 1st differencing, adding a constant term to these estimates, and correcting for autocorrelation where needed, models 0 and 2-4 had 60 degrees of freedom and model 1 had 44.

Table 2: Translog Production Function Empirical Results*

	0. No correction	1. Obs. with maintenance deleted	2. Net capital	3. Net capital and net labor	4. Adjusted output & lignite
I. Elasticities					
ϵ lignite	0.776	0.550	0.767	0.811	0.904
ϵ capital	0.068	0.168	0.069	0.062	0.010
ϵ labor	0.093	0.230	0.096	0.053	0.014
RTS	0.937	0.948	0.932	0.926	0.928
II. T-statistics					
Ln(lignite)	-2.071	-1.920	-2.215	-2.140	-0.871
Ln(capital)	1.547	0.702	1.501	2.451	2.098
Ln(labor)	1.270	2.317	1.273	0.566	-0.526
$\frac{1}{2}(\text{Ln}(\text{lignite}))^2$	-1.462	-1.289	-1.394	-2.324	-0.433
$\frac{1}{2}(\text{Ln}(\text{capital}))^2$	-3.775	-1.680	-2.949	-3.731	-4.069
$\frac{1}{2}(\text{Ln}(\text{labor}))^2$	-2.687	-2.318	-3.092	-1.781	-0.781
Ln(capital)*Ln(labor)	1.480	1.120	1.208	0.839	1.024
Ln(capital)*Ln(lignite)	0.359	-0.239	1.070	1.203	0.921
Ln(labor)*Ln(lignite)	2.630	2.420	2.313	2.207	0.305
III. R-Bar²					
	0.999	0.999	0.999	0.999	0.999

* ϵ = marginal product output elasticity and RTS = returns to scale. The data was first differenced because Dickey-Fuller tests indicated that 1st differencing was the correct way of detrending the data (instead of adding a trend term). The Durbin-Watson statistics indicated that all production net of maintenance models needed to be corrected for first degree autocorrelation (and that a first degree correction was sufficient). After first differencing, adding a constant term, and correcting for autocorrelation, models 0 and 2-4 had 52 degrees of freedom and model 1 had 37 degrees of freedom. Since the translog production function generates a different elasticity and returns to scale for each observation, the mean elasticity and returns to scale are reported here.

Most signs in the regressions which underlie Tables 1 and 2 are as expected, with two notable exceptions. Under the Cobb-Douglas functional form, the sign on the elasticity of capital was perverse and under the Translog functional form, the sign on $\ln(\text{lignite})$ was perverse (but the elasticity of lignite was not perverse). These perverse signs could be due to pollution abatement efforts. Capital and lignite are collinear because capital had to be expanded so that more lignite could be burned. As capital was expanded and more lignite burned, more sulfur dioxide was released which resulted in more acid rain during winter temperature inversions (Leightner and Lovell, 1998 and Leightner, 1999). Thus each winter (in comparison to the rest of the year), Mae Moh used a smaller percent of its capacity and burned less lignite per generating unit due to its pollution abatement efforts. These pollution abatement efforts could explain why the sign on lignite (under Translog) and the sign on capital (under Cobb-Douglas) were perverse – the collinear relation between capital and lignite would cause either capital or lignite to capture the effect of both on output and the other to capture the effects of increased pollution abatement efforts.

A close examination of Panel A of Tables 1 and 2 reveals a reoccurring pattern which produces the following rank order of models (see Section III for a description of these models):

Model	ϵ labor	ϵ lignite
1. Maintenance observations deleted	highest	lowest
2. Net capital		
0. No correction	middle	middle
3. Net capital and net labor		
4. Adjusted output and adjusted lignite	lowest	highest

When compared to no correction for maintenance (model 0), I had hoped that all methods of correcting for maintenance would produce similar results.

Contrary to my hope, models 1 and 2 increased the elasticity of labor and decreased the elasticity of lignite, while models 3 and 4 did the opposite. Furthermore, the magnitude of the changes in elasticities are large. When compared to no correction for maintenance (model 0) and when a Cobb-Douglas production function was used, the elasticity of lignite fell by 4.5 percent when all maintenance months were deleted, but rose by 2.6 percent when output and lignite were adjusted; the elasticity of capital fell by 92.1 percent when all maintenance months were deleted, but rose by 18.4 percent when output and lignite were adjusted; the elasticity of labor rose by 19.8 percent when all maintenance months were deleted, but fell by 25.5 percent when output and lignite were adjusted. Furthermore, in every case the Translog production function produced even larger variations in estimated elasticities. Not correcting for maintenance and the second best method used for correcting for maintenance makes a large difference in the estimated elasticities of all inputs, not just the input undergoing maintenance.

Remember that these output elasticities are directly related to the marginal products of the inputs and that profit maximizing rules for the hiring of inputs are based on these marginal products. Therefore, variations in output elasticities like those shown in Panel A of Tables 1 and 2 would drastically affect a firm's calculation of its optimal combination of inputs and, thus, its profits.

Although the output elasticities which emerge from the four estimated models vary drastically, the measures of returns to scale (RTS) vary much less. The RTS measure of 0.961 for the Cobb-Douglas production function estimate when all maintenance months are eliminated implies that if all inputs are increased by 100%, then output would increase by 96.1%. All the other RTS measures for the Cobb-Douglas production function technique imply that doubling all inputs would more than double output. In contrast, the translog production function technique always generated returns to scale measures that indicate that doubling all inputs would fall short of doubling output.

How production net of maintenance is modeled and the technique used drastically affect the output elasticities and the measures of returns to scale

which emerge from production function estimations. By so doing, maintenance affects the calculated optimal combination of inputs in the short run and the optimal long run level of all inputs.

V. Conclusion

Maintenance and production interact. Maintenance either produces an immediate increase in the productivity of the input maintained or it increases the longevity of the input. Both of these types of maintenance result in the input producing more output over its entire life. Researchers estimating production or cost functions should take the interaction between production and maintenance into account by adjusting the temporal length of their observations. If maintenance is longevity producing, then (in an ideal world) the temporal length of observations should equal the life span of the input. If maintenance is marginal product increasing, then the temporal length of observations should be integer multiples of production-maintenance cycles.

Often this ideal method of correcting for maintenance is impossible for large firms. In this paper I tested four different second best methods of accounting for maintenance when estimating Cobb-Douglas and Translog production functions. These four models produced very different output elasticities for the traditional inputs. When compared to no correction, two of these models increased the output elasticity of labor and decreased the output elasticity of lignite. The other two models did the opposite. Hopefully, future research will indicate which of these second best models is best.

Finally, maintenance raises several different issues which are only tangentially related to this paper's analysis. First, what is the optimal amount of maintenance? Second, how can maintenance saving technology be modeled and what are its effects on the marginal products of other inputs and on the longevity of the input normally maintained? Third, what are the effects of routine versus emergency maintenance on a firm's activities and objectives? Finally, what are the possible gains from specializing inputs for the conducting

of maintenance and what forces limit these gains from specialization? Hopefully this paper will initiate an extensive study of how maintenance affects production estimations, which we have ignored in the past.

References

- Jorgenson, D.W. (1996), "Empirical Studies of Depreciation," *Economic Inquiry*, 34: 24-42.
- Leightner, J. and Lovell, C.A.K. (1998), "Weather and Pollution Abatement Costs," *The Energy Journal*, 19: 2, 165-189.
- Leightner, J. (1999), "Weather-Induced Changes in the Trade off between SO₂ and NO_x at Large Power Plants," *Energy Economics*, 21: 239-259.
- Rothwell, Geoffrey and Rust, John. "A Dynamic Programming Model of U.S. Nuclear Power Plant Operations," Published on the World Wide Web by Econ WPA at <http://econwpa.wustl.edu/eprints/mic/papers/9502/9502001.abs>.
- Triplett, J.E. (1996), "Depreciation in Production Analysis and in Income and Wealth Accounts: Resolution of an Old Debate," *Economic Inquiry*, 34: 93-115.