



91740127

Relais Request No. URG-23129531

Customer Code

89-0096

Delivery Method

SED

Request Number

RZ148422 SED99*2*

Scan

Date Printed: 19-Jan-2007 11:58

Date Submitted: 19-Jan-2007 11:15

4942.535000

TITLE: Journal of applied economics -Bu...

YEAR: 1994

VOLUME/PART: 1994 VOL 9 PT 4 PP 351-368

PAGES:

AUTHOR:

ARTICLE TITLE:

SHELFMARK: 4942.535000

Your Ref :

RZ148422 SED99*2*|JOURNAL OF APPLIED ECONOMICS|1994 VOL 9 PT 4 PP 351-368|D
MARK KENNET|A STRUCTURAL MODEL OF AIRCRAFT...|BLIN

DELIVERING THE WORLD'S KNOWLEDGE**This document has been supplied by the British Library****www.bl.uk**

The contents of the attached document are copyright works. Unless you have the permission of the copyright owner, the Copyright Licensing Agency Ltd or another authorised licensing body, you may not copy, store in any electronic medium or otherwise reproduce or resell any of the content, even for internal purposes, except as may be allowed by law.

The document has been supplied under our Copyright Fee Paid service. You are therefore agreeing to the terms of supply for our Copyright Fee Paid service, available at :

www.bl.uk/services/document/edd.html

A STRUCTURAL MODEL OF AIRCRAFT ENGINE MAINTENANCE

D. MARK KENNET

*Office of Economic Research, US Bureau of Labor Statistics, Suite 4915, 2 Massachusetts Avenue, NE,
Washington, DC 20212, USA*

SUMMARY

We develop and estimate a simple regenerative optimal stopping model of aircraft engine maintenance that attempts to describe the behaviour of airline maintenance personnel. The model assumes that the decision to send an engine to the shop for overhaul is the solution to a stochastic dynamic programming problem that trades off the expected cost of continuing operation with the attendant risk of engine failure with the cost of performing the overhaul. We estimate the model using 42 engine histories from Pratt & Whitney, Inc. Estimation results indicate that such a model does not explain observed engine histories before deregulation, but does fit the data in the era since deregulation. The model also provides insight into the perceived relative costs of engine maintenance, in-flight shutdown, and ordinary operation.

1. INTRODUCTION

This paper develops a dynamic model of aircraft engine maintenance based on Rust's (1987) bus maintenance model. Our intention is to determine the structural changes that led to the result that deregulation enacted by Congress in 1978 influenced airlines to fly aircraft engines longer before overhauling them (see Kennet, 1993).¹ In the earlier work, we employed a *reduced-form* duration model to conclude that 42 engine histories collected from Pratt & Whitney, Inc. exhibit evidence that airlines' behaviour has changed.

This study exploits the notion that efficient firms either before or after deregulation would choose a maintenance policy that minimizes discounted expected cost by solving a stochastic discrete dynamic programming (SDDP) problem. SDDP determines the probability of observing any particular action on the part of the firm, and thus the maximum likelihood estimation used to estimate the model must incorporate the solution to SDDP in the computation of the likelihood function. Thus, we provide further experience in estimating 'bottom-up' models like those of Rust (1987), Das (1992), and Sturm (1990).

We will argue that the change in firm behaviour from the regulated to unregulated era results from a change in the firms' perception of the relative weights of the cost of having an engine visit the shop versus the cost of continuing operation (with its attendant risk of engine failure).

¹ See also Rose (1989, 1990, 1992) for results establishing that overall safety has not been affected by the fluctuating profits resulting from deregulation. Rose (1992) surveys other literature analysing National Transportation Safety Board accident data as well as passenger fatality rates.

2. BACKGROUND

Although aircraft maintenance policies are essentially condition-based, inspection (i.e. the means of determining the system's condition) is quite costly.² Sherwin (1979) derives an optimal inspection interval for one of the two major maintenance policies. However, inspection intervals are more or less exogenously determined for airlines through regulations administered by the Federal Aviation Administration (FAA).

Based on the inspections of engines, airline maintenance engineers determine the point at which an engine requires a visit to the shop for an overhaul. Thus, the data used in this study are complete engine histories for 42 Pratt & Whitney engines. The histories include the dates and number of operating hours on the engine for all 'events' in the engines' lifetimes, which begin as early as 1964 and end in 1988. Summary statistics and reduced-form analysis are found in Kennet (1993).

Duration results in Kennet (1993) suggest that jet engines are likely to exhibit increasing failure rates; that is, the longer the engine has been operating since its last overhaul, the more likely it is to fail. Sherif and Smith (1981) point out that operations research results indicate that progressive maintenance scheduling is optimal in this situation.³ The continual updating of the information and response to that information attendant to progressive maintenance leads us to the model in the next section, which owes a large debt to Rust (1987).

3. THE MODEL

We assume that airline j at calendar time t chooses

$$d_{jt} = \begin{cases} 0 & \text{do nothing} \\ 1 & \text{remove engine for shop visit} \end{cases} \quad (1)$$

These choices depend on the state variable x_{jt} , which is observable to both maintenance personnel and the econometrician. In this study, x_{jt} is two-dimensional and includes hours of operation of the engine and an indicator of whether the engine has experienced a shutdown since the last shop visit (henceforth, we will suppress j for simplicity).

Also influencing the decision process is ε_t , a vector of engine state variables observable to airline maintenance personnel but not to the econometrician. The elements of ε_t might be functions of boroscope results, exhaust gas temperature readings, and engine pressure ratio readings, as well as intensity indicators such as engine 'cycles', or hours of operation at full throttle.⁴ ε_t has two elements, corresponding to each of the decision values. Each element is interpreted as a component of utility for the airline associated with the decision d . That utility is assumed to be $u(x_t, d, \theta_1) + \varepsilon_t(d)$, where θ_1 is a vector of unknown parameters.

² According to Edward R. Cowles, Public Relations Director for Pratt & Whitney, Inc., there are generally two types of maintenance strategies: 'on-condition' and 'condition monitoring'. On-condition maintenance depends on inspections and tests, which also give the locations for work to be performed. Condition monitoring involves allowing noncritical parts to fail and using the failure as a signal for maintenance.

³ The Navy defines a progressive policy as 'adjusting rework content and frequency as necessary to preclude the need for overhaul, and to assure, within high confidence limits, continuance of a material condition which will sustain the aircraft through a subsequent operating tour,' (see Schwartz *et al.*, (1971). Essentially, progressive maintenance is a flexible policy (as opposed to a rigid policy of always performing work after a part has been used a certain number of hours) in which minor discrepancies uncovered during an inspection are overlooked if they do not affect safety, thus putting off the expense of an overhaul. The engine is overhauled only when normal maintenance procedures can no longer ensure safe operation.

⁴ Engine cycles are included in the raw data we obtained from Pratt & Whitney, but there seem to be so many missing values and obviously erroneous values that we must treat cycles as an unobservable.

Let $\phi(x_{t+1}, \varepsilon_{t+1} | x_t, \varepsilon_t, d_t, \theta_2, \theta_3)$ be the Markov transition density for the state variables (x_t, ε_t) associated with choice d_t and given the unknown parameter θ_2 and θ_3 . Finally, define $\theta := (\beta, \theta_1, \theta_2, \theta_3)$ as the complete vector of parameters to be estimated, with β as the discount factor. The dimensions of θ is $1 + K_1 + K_2 + K_3$, where K_1 , K_2 , and K_3 are the dimension of θ_1 , θ_2 and θ_3 , respectively.

One goal in this study is to determine whether these data reveal a stationary decision rule for all periods; that is, whether or not deregulation brought about a change in the decision deemed optimal before deregulation. To discuss the optimal rule, we require the optimal value function that is the solution to Bellman's equation:

$$V_\theta(x_t, \varepsilon_t) = \max_{d \in C(x_t)} [u(x_t, d, \theta_1) + \varepsilon_t(d) + \beta EV_\theta(x_t, \varepsilon, d)] \quad (2)$$

where

$$EV_\theta(x_t, \varepsilon_t, d) := \int_y \int_\tau V_\theta(y, \tau) \phi(dy, d\tau | x_t, \varepsilon_t, d, \theta_2, \theta_3) \quad (3)$$

A stationary decision rule is defined as

$$d_t = f(x_t, \varepsilon_t, \theta) \quad (4)$$

where

$$f(x_t, \varepsilon_t, \theta) := \operatorname{argmax}_{d \in C(x_t)} [u(x_t, d, \theta_1) + \varepsilon_t(d) + \beta EV_\theta(x_t, \varepsilon_t, d)] \quad (5)$$

is the optimal control.

It is possible to estimate equation (4) by numerically integrating over equation (3) to obtain a conditional choice probability $P(d_t | x_t, \theta)$. However, Rust (1987) points out that the conditional independence assumption, namely, that

$$\phi(x_{t+1}, \varepsilon_{t+1} | x_t, \varepsilon_t, d, \theta_2, \theta_3) = q(\varepsilon_{t+1} | x_{t+1}, \theta_2) p(x_{t+1} | x_t, d, \theta_3) \quad (6)$$

reduce the computational burden. The conditional independence assumption also reduces the state space for computing the fixed point EV_θ to the set $\{(x, d) | x \in \mathbb{R}^K, d \in C(x)\}$ rather than the larger set of all combinations of x and the unobservable ε . With this simplification, it is possible to express the log likelihood for an observation as

$$l_t = \ln(P(d_t | x_t, \theta)) + \ln(p(x_t | x_{t-1}, d_{t-1}, \theta_3)) \quad (7)$$

that is, the log of the joint distribution function of d_t and x_t , and to sum over t for the log likelihood for the sample.

We will now specify the functional forms more precisely. Utility is defined as the negative of costs for a given alternative plus the value of the unobservable $\varepsilon_t(d)$ as described earlier. Cost is defined as engine removal and repair cost, SVC ('shop visit cost'), if any, plus the variable cost associated with the state of the system, $c(x_t, \theta_1)$. Thus, we have

$$u(x_t, d, \theta_1) + \varepsilon_t(d) = \begin{cases} -c(x_t^1, x_t^2, \theta_1) + \varepsilon_t(0) & \text{if } d = 0 \\ -SVC - c(0, 0, \theta_1) + \varepsilon_t(1) & \text{if } d = 1 \end{cases} \quad (8)$$

Note that if $d = 1$ in the model, then $x_t = (0, 0)$ (the observed two-dimensional state variable reverts to the zero state). The transition probability for x_t is bivariate, and adopting the convention that the first element of x is the hours of operation and the second element the engine shutdown history indicator:

$$p(x_{t+1} | x_t, d_t, \theta_3) = g(x_{t+1}^2 | x_{t+1}^1, x_t, d_t, \theta_3) \cdot h(x_{t+1}^1 | x_t, d_t, \theta_3) \quad (9)$$

We assume that $h(x_{t+1}^1 | x_t, \theta_3) = h(x_{t+1}^1 | x_t^1, \theta_3)$; that is, that previous shutdown information has no effect on current hours of operation.⁵ Given this assumption, we can discretised monthly hours into two groups of a given length.

The distribution g is somewhat more difficult. We assume that

$$\Pr\{x_{t+1}^2 = 1 | x_t^2 = 1\} = 1 \quad (10)$$

and that g is nondecreasing in x_{t+1}^1 . For the cases in which x_t^2 is 0 (i.e. there is no current shutdown history) it may be argued that the probability of a shutdown occurring in an interval is not independent of where in history that interval lies. However, Kennet (1993) finds that hours since shop visit is not statistically significant in predicting an engine shutdown. Thus, we claim justification in assuming that the density function g depends only on x_t^2 and the interval $x_{t+1}^1 - x_t^1$ and that the transition probabilities are multinomial.⁶

With this extra assumption, we can write down the Markov probability transition matrix after we discretize engine hour data appropriately. If 35,000 hours is the upper bound on time between renewals (this is a value greater than any observed in the data), the discretized state variable corresponding to hours is computed by dividing 35,000 by the assumed dimension of the value function. For a 44-dimensional value function, each hour state x_t^1 corresponds to about 795 hours.

Although the state variable x_t is two-dimensional, it is possible to express the case in a one-dimensional vector by pairing the hours component with a zero for no shutdowns and a one for shutdowns. For example, state 0 is engine hours between 0 and 795, no shutdowns; state 1 is engine hours between 0 and 795, one or more shutdowns; state 2 is engine hours greater than 795 but less than 1590, no shutdowns; state 3 is engine hour greater than 795 but less than 1590, one or more shutdowns. Redefining the θ_3 vector, the 88-dimensional transition matrix for the uncontrolled Markov process (that is, with no shop visits) is

	0	1	2	3	4	5	...	84	85	86	87
0	θ_{30}	θ_{31}	θ_{32}	θ_{33}	0	0	...	0	0	0	0
1	0	θ_{34}	0	θ_{35}	0	0	...	0	0	0	0
2	0	0	θ_{30}	θ_{31}	θ_{32}	θ_{33}	...	0	0	0	0
3	0	0	0	θ_{34}	0	θ_{35}	...	0	0	0	0
4	0	0	0	0	θ_{30}	θ_{31}	...	0	0	0	0
$\pi = 5$	0	0	0	0	0	θ_{34}	...	0	0	0	0
\vdots	\vdots	\vdots	\vdots	\vdots	\vdots	\vdots	\vdots	\vdots	\vdots	\vdots	\vdots
84	0	0	0	0	0	0	...	θ_{30}	θ_{31}	θ_{32}	θ_{33}
85	0	0	0	0	0	0	...	0	θ_{34}	0	θ_{35}
86	0	0	0	0	0	0	...	0	0	θ_{30}	$1 - \theta_{30}$
87	0	0	0	0	0	0	...	0	0	0	1

(11)

Note the difference between even rows and odd rows of π . The odd rows are parameterized differently because of the construction of the shutdown component of the state variable. Once an engine experiences a shutdown in its shop visit cycle, it cannot revert to the no-shutdown state. This construction reflects an indication in the raw data that a *history* of shutdowns—

⁵ Kennet (1993) finds that engine duration is unaffected by shutdown history.

⁶ We also estimated a version of the model in which this assumption is relaxed and the probability of shutdown is allowed to vary with the engine history. Results were qualitatively similar.

rather than a shutdown *per se*—leads to a decision to send an engine to the shop. (Odd rows correspond to the 'one or more shutdown state'.) Thus, the probability of going from an odd to an even state is always zero. The probability θ_{34} , is assumed to differ from θ_{30} because we wish to allow the probability of transition between hour-states when a shutdown has occurred to differ from that when a shutdown has not occurred. As we shall see later, the data support this specification, although Kennet (1993) finds that shutdown history does not predict hours of operation between shop visits in a reduced-form framework.

It is also worth pointing out that although six parameters are specified in equation (11), only four require estimation because $\theta_{33} = 1 - \theta_{30} - \theta_{31} - \theta_{32}$ and $\theta_{35} = 1 - \theta_{34}$. Also note that state 87, by assumption, is 35,000 hours or more; thus, state 87 is an absorbing state, and cell (87, 87) is exactly unity.

The cost function is most conveniently expressed in terms of the original state variable rather than the vectorized form. We propose

$$c(x_t, \theta_t) = 0.001 * \theta_{11} * \text{hours}_t + \theta_{12} * \text{shutdown}_t \quad (12)$$

Note that shutdown_{*t*} is simply a dummy variable defined for a shutdown history. This function simply adds a fixed shutdown cost to the operation cost.

The assumption of bivariate extreme value distribution for ε_t implies that $P(d_t | x_t, \theta)$ is logistic, with the logit probability for each decision depending on the solution of equation (4):

$$P(d_t = \text{shop visit} | x_t, \theta) = \frac{\exp(-SVC - c(0, 0, \theta) + \beta EV_\theta(0))}{1 + \exp(-SVC - c(0, 0, \theta) + \beta EV_\theta(0) + c(x_t, \theta) - \beta EV_\theta(x_t))} \quad (13)$$

and

$$P(d_t = \text{do nothing} | x_t, \theta) = \frac{1}{1 + \exp(-SVC - c(0, 0, \theta) + \beta EV_\theta(0) + c(x_t, \theta) - \beta EV_\theta(x_t))} \quad (14)$$

Thus, to solve the implied maximum likelihood problem, we employ a nested fixed point algorithm, which 'nests' a Newton-type iterative procedure for evaluating the likelihood function inside the likelihood maximizing algorithm. See Rust (1987, 1988) for mathematical details.

The practicality of estimating the parameters of this model is enhanced by dividing equation (7) into two segments, $P(d_t | x_t, \theta)$ (the choice probability) and $p(x_t | x_{t-1}, d_{t-1}, \theta_3)$ (the state variable transition probabilities). In a two-stage procedure, the choice probability part is estimated first, using nonparametric estimates of the transition probabilities. In the second stage, the full maximum likelihood estimation is performed, coupling the estimated parameters from the partial estimation with the nonparametric transition probability estimates as starting values.

4. RESULTS AND ANALYSIS

In performing the maximum likelihood estimation the maximization algorithms drove the value of β to unity.⁷ Thus, in reporting the result of this study, we will maintain the convention of reporting estimates with β held fixed at both zero and 0.99999. We are also

⁷ We used the Berndt, Hall, Hausman, and Hall (BHHH) algorithm for most of the work performed in this study, although we tried the Broyden, Fletcher, Goldfarb, and Shanno (BFGS) and the Nelder-Mead downhill simplex methods to determine whether other optimization routines corrected this problem. They did not. See Press *et al.* (1986) for details of the latter two algorithms.

reporting results for β equal to 0.9923007, which corresponds to the weighted average cost of capital of eight representative airlines.⁸

Table I reports full likelihood results for JT8D engines under the regulated and unregulated regimes with $\beta = 0$, $\beta = 0.9923007$, and $\beta = 0.99999$. Table II reports analogous results for JT9D engines.⁹ We assume that the cost function is linear in the two state variables as in equation (13).¹⁰ Thus, operating cost in the tables corresponds to θ_{11} and shutdown cost to θ_{12} .

Because the parameters in this model are only identified up to scale, we report normalized values to permit heuristic comparisons. The normalized values are simply the actual estimate divided by the \mathbb{L}^1 norm of the cost parameter vector. Standard errors reported are those of the normalized values, computed using the δ -method (see Billingsley, 1979). Second, we used a logistic transformation designed to guarantee both that the transition probability matrix parameters lie in the (0, 1) interval and that the matrix rows sum to unity. The estimation algorithm reports the transformed probabilities and the standard errors of the transformations (these are the Q parameters in the tables). We also report actual estimated probabilities underneath the standard errors.

The heterogeneity tests in the tables refer to a likelihood ratio test of whether the data support separating the sample into regulated and unregulated eras. As the tables demonstrate, the test statistic is significant, indicating that firms' decision rules were different in the two eras.

Another potential source of heterogeneity is from the firms themselves. That is, firms may follow different decision rules when faced with the same value of the state variable. To test this notion, we estimated the model separately for each firm for each engine type under each regime. The regulated years sample for both engine types yielded an insignificant likelihood ratio statistic, supporting the null hypothesis that no firm effects were present. The deregulated years sample rejected the null under at least two of the assumed values for β . Table III summarizes the results for JT8D engines and Table IV for JT9D engines. Table V summarizes the likelihood ratio statistics.

For purposes of comparison we have included Table VI, which contains the results of a binary logit estimation using the occurrence of a shop visit as the dependent variable and the two state variables, discretized hours and shutdown history, as explanatory variables. Note first that these reduced-form results for the JT8D case make no sense based on any reasonable expectation of a relationship between hours of operation and the likelihood of shop visit: The coefficient corresponding to the structural operating cost (θ_{11}) is negative. Second, observe that there is no possibility of incorporating into this reduced-form 'model' any notion of optimization over time. We cannot even report different results for differing β 's. Finally, note

⁸ The eight airlines are USAir, American, Continental, Eastern, Northwest, Pan Am, TWA, and United. Ideally, we would have employed the individual firms' weighted average cost of capital. However, we are not permitted to know the identity of the firms owning the engines in the sample. We computed the cost of capital by weighting firms' cost of equity (based on Value Line's estimate of firm CAPM β s) and cost of debt (based on long-term bond rates from Salomon Brothers, Inc.) for the sample years available. Weights were determined by firms' debt-to-equity ratios from Value Line. The average annual weighted cost of capital was 0.0931085, which was divided by 12 to arrive at a monthly weighted cost of capital equal to 0.0077590. The discount factor, β , was then computed as $1/(1 + r)$, where $r = 0.007590$.

⁹ JT8D engines have approximately 16,000 lb of thrust and are generally used on Boeing 727 and McDonnell-Douglas DC-9 aircraft. JT9D engines have 41,000 lb of thrust and are generally used on Boeing 747s.

¹⁰ We attempted to specify alternative cost structures in order to ascertain the validity of the linear model. A model quadratic in hours of operation failed to converge in the JT8D engine cases. As an alternative, we estimated a version in which costs are quadratic in the engine-hour state variable. We also estimated a partial likelihood version in which we relax the assumption that the probability of shutdown depends only on the interval $x_t - x_{t-1}$. The respecification is to vary θ_{31} and θ_{33} such that they take on one value if hours are less than 2385 and a different value if hours are greater than 2385. Neither of the above respecifications produces qualitatively different results.

Table I. Structural parameter estimates for JT8D engines (linear cost)

Parameter	Regulated era			Unregulated era		
	0	0.9923	0.9999	0	0.9923	0.9999
<i>Cost parameters</i>						
Shop visit cost	0.756 (0.027)	0.958 (0.036)	0.960 (0.036)	0.728 (0.032)	0.956 (0.050)	0.959 (0.050)
Operating cost	0.000 (0.007)	2e-04 (0.001)	0.000 (0.001)	0.015 (0.005)	0.002 (0.001)	0.002 (0.001)
Shutdown cost cost	0.243 (0.048)	0.042 (0.010)	0.040 (0.010)	0.257 (0.086)	0.042 (0.020)	0.039 (0.019)
<i>Transition matrix parameters</i>						
Q_0	5.435 (0.334)	5.463 (0.335)	5.464 (0.335)	6.649 (0.709)	6.717 (0.711)	6.719 (0.710)
θ_{30}	0.768	0.768	0.768	0.737	0.738	0.738
Q_1	1.610 (0.365)	1.609 (0.365)	1.609 (0.365)	1.219 (0.802)	1.253 (0.802)	1.253 (0.802)
θ_{31}	0.017	0.016	0.016	0.003	0.003	0.003
Q_2	4.150 (0.336)	4.181 (0.336)	4.181 (0.337)	5.601 (0.710)	5.668 (0.711)	5.670 (0.711)
θ_{32}	0.212	0.213	0.213	0.258	0.258	0.258
Q_4	-1.182 (0.131)	-1.209 (0.131)	-1.209 (0.131)	-1.972 (0.340)	-1.042 (0.340)	-1.043 (0.340)
θ_{34}	0.765	0.770	0.770	0.726	0.739	0.739
- Log likelihood	2776.4	2776.3	2776.3	1839.4	1837.6	1837.6
Degrees of freedom		3094			2280	
Myopia test: LR test statistic		0.0	0.2		3.6	3.6
Significance (df = 1)		0.99	0.89		0.06	0.06
Heterogeneity test: LRT			61.4			
Significance (df = 1)			0.0			

Standard errors in parentheses.
Values reported are normalized.

that while we cannot make a direct comparison between this reduced-form specification and the structural model developed in this paper, the partial likelihood formulation of the structural model has the same number of parameters and hence the same number of degrees of freedom as the reduced form, and so we advisedly report that for all four regimes under discussion here (JT8D regulated, JT8D unregulated, JT9D regulated, JT9D unregulated) except one, the log likelihood for the sample in the reduced-form cases was larger than the log likelihood for the sample in the structural cases, and in two of the four a likelihood ratio test would reject the null hypothesis that the likelihood of observing the sample is less under the

Table II. Structural parameter estimates for JT9D engines (linear cost)

Parameter	Regulated era			Unregulated era		
	0	0.9923	0.9999	0	0.9923	0.9999
<i>Cost parameters</i>						
Shop visit cost	0.709 (0.033)	0.905 (0.046)	0.907 (0.046)	0.774 (0.030)	0.961 (0.050)	0.962 (0.051)
Operating cost	0.088 (0.018)	0.027 (0.007)	0.027 (0.007)	0.023 (0.005)	0.003 (0.001)	0.003 (0.001)
Shutdown cost	0.203 (0.051)	0.067 (0.019)	0.055 (0.019)	0.203 (0.042)	0.036 (0.009)	0.034 (0.009)
<i>Transition matrix parameters</i>						
Q_0	4.226 (0.262)	4.278 (0.262)	4.278 (0.262)	4.244 (0.239)	4.242 (0.238)	4.247 (0.238)
θ_{30}	0.738	0.739	0.739	0.655	0.653	0.653
Q_1	1.512 (0.285)	1.555 (0.285)	1.555 (0.285)	0.871 (0.282)	0.848 (0.282)	0.848 (0.282)
θ_{31}	0.049	0.048	0.048	0.022	0.022	0.022
Q_2	2.929 (0.267)	2.983 (0.267)	2.982 (0.267)	3.505 (0.241)	3.515 (0.240)	3.519 (0.240)
θ_{32}	0.202	0.202	0.202	0.313	0.315	0.315
Q_4	-1.020 (0.124)	-1.020 (0.124)	-1.020 (0.124)	-0.446 (0.104)	-0.426 (0.104)	-0.425 (0.104)
θ_{34}	0.735	0.735	0.735	0.610	0.605	0.605
- Log likelihood	2089.5	2089.5	2089.5	2384.3	2379.5	2379.4
Degrees of freedom		1792			2295	
Myopia test: LR test statistic		0.1	0.0		9.6	9.8
Significance (df = 1)		0.75	0.83		0.00	0.00
Heterogeneity test: LTT				184.8		
Significance (df = 1)				0.0		

Standard errors in parentheses.
Values reported are normalized.

structural specification than under the reduced form. Thus, the benefit of developing a structural model seems clear.

A fairly striking result is the failure of what Rust (1987) calls the 'myopia test' in the regulated era data and its passage in the unregulated era. This statistic essentially provides a test of the significance of β , the discount factor, although it is not possible to draw strong inferences because β is not truly identified. An implication of this result—that estimation with β zero or positive cannot be distinguished during the regulated era but can be distinguished during the unregulated era—is that maintenance during the regulated years was performed at a superoptimal level (i.e. as we might gather from the fact that shop visits occurred more

Table III. Firm-specific estimation results deregulated years for JT8D engines

β	0.0000			0.9923			0.9999		
	Shop visit cost	Operating cost	Shut-down cost	Shop visit cost	Operating cost	Shut-down cost	Shop visit cost	Operating cost	Shut-down cost
Firm									
1	0.708 (0.041)	0.028 (0.008)	0.264 (0.093)	0.952 (0.062)	0.004 (0.001)	0.044 (0.022)	0.854 (0.063)	0.003 (0.001)	0.042 (0.022)
2	0.746 (0.066)	0.014 (2e-04)	0.240 (0.212)	0.590 (0.069)	0.001 (2e-04)	0.409 (0.157)	0.766 (0.085)	0.001 (2e-04)	0.233 (0.205)
3	0.708 (0.067)	0.016 (2e-04)	0.275 (0.243)	0.648 (0.072)	0.002 (2e-04)	0.349 (0.204)	0.725 (0.080)	0.002 (2e-04)	0.273 (0.240)
4	0.743 (0.087)	0.020 (2e-04)	0.236 (0.153)	0.972 (0.105)	0.007 (1e-04)	0.021 (0.042)	0.974 (0.102)	0.007 (1e-04)	0.019 (0.030)
5	0.799 (0.070)	0.016 (2e-04)	0.186 (0.155)	0.199 (0.032)	1e-04 (4e-05)	0.800 (0.046)	0.779 (0.167)	4e-04 (2e-04)	0.221 (0.172)
6	0.934 (0.125)	0.016 (2e-04)	0.050 (0.244)	0.914 (0.079)	0.004 (8e-05)	-0.802 (0.071)	0.923 (0.081)	0.004 (8e-05)	-0.073 (0.067)
7	0.712 (0.101)	0.016 (2e-04)	0.272 (0.240)	0.646 (0.117)	0.004 (2e-04)	-0.350 (0.088)	0.910 (0.146)	0.010 (2e-04)	0.080 (0.160)

Table IV. Firm-specific estimation results deregulated years for JT9D engines

β	0.0000			0.9923			0.9999		
	Shop visit cost	Operating cost	Shut-down cost	Shop visit cost	Operating cost	Shut-down cost	Shop visit cost	Operating cost	Shut-down cost
Firm									
1	0.744 (0.039)	0.004 (0.007)	0.252 (0.054)	0.952 (0.057)	0.002 (0.001)	0.046 (0.012)	0.954 (0.057)	0.002 (0.001)	0.044 (0.011)
2	0.747 (0.100)	0.076 (0.026)	0.197 (0.121)	0.937 (0.150)	0.016 (0.007)	0.047 (0.035)	0.938 (0.150)	0.016 (0.007)	0.046 (0.035)
3	0.0889 (0.133)	0.097 (0.030)	0.014 (0.135)	0.977 (0.163)	0.018 (0.007)	-0.005 (0.034)	0.977 (0.164)	0.018 (0.006)	-0.005 (0.033)
4	0.707 (0.124)	0.036 (0.020)	-0.258 (0.232)	0.943 (0.254)	0.007 (0.005)	-0.050 (0.041)	0.945 (0.258)	0.006 (0.005)	-0.048 (0.039)

frequently before deregulation, maintenance was performed too often), and the pressures of competition under deregulation caused airlines to cut costs (and shop visits) to an optimal level. This result alone could explain the significant coefficient on deregulation in the analysis in Kennet (1993).

This result also gains credence from the fact that only in the deregulated years was the firm heterogeneity significant. A plausible explanation is that prior to deregulation, firms followed similar maintenance rules: a large percentage of airline maintenance personnel received their initial training in the military, and there may have been a widespread 'by the book' approach

Table V. Firm heterogeneity test results

β Regime	0.0000			0.9923			0.999		
	LR test statistic	Degrees of freedom	Significance	LR test statistic	Degrees of freedom	Significance	LR test statistic	Degrees of freedom	Significance
JT8D regulated	37.06	42	0.6873	17.08	42	0.9998	37.94	42	0.6498
JT8D deregulated	46.50	42	0.2923	54.87	42	0.0880	56.00	42	0.0727
JT9D regulated	26.86	21	0.1755	27.78	21	0.1465	27.84	21	0.1447
JT9D deregulated	45.28	21	0.0016	40.81	21	0.0059	40.72	21	0.0061

Table VI. Illustrative reduced form results binary logit

Parameter	JT8D regulated	JT8D unregulated	JT9D regulated	JT9D unregulated
Intercept	-2.500 (0.109)	-3.246 (0.138)	-2.195 (0.128)	-2.782 (0.125)
Operating cost	-0.006 (0.026)	0.064 (0.017)	0.247 (0.052)	0.048 (0.022)
Shutdown cost	0.862 (0.162)	1.306 (0.378)	0.545 (0.138)	0.706 (0.172)
- Log likelihood	886.2	491.7	811.6	649.0
Degrees of freedom	3098	2284	1796	2299
LK statistic*	0.1	3.2	0.0	5.5

* Likelihood ratio statistic 'testing' the null that the samples are not more likely under the structural (partial likelihood with $\beta = 0.9999$) versus the reduced-form specifications.

Note: - Log likelihood values or reduced-form logit specifications were greater than or equal to those for the structural specifications in all cases. Thus, an interpretation of these values is that (at least during the unregulated era) the structural model fit the data better than the reduced-form model.

to maintenance by all airlines. After deregulation, firms began refining and optimizing their maintenance scheduling; hence, both the firm heterogeneity and the better fit in general of the SDDP model.

This story has some theoretical backing; for example, Kamien and Vincent (1990) present a model that supports earlier claims by Panzar and Savage (1987), Stiglitz and Arnott (1987), and Braeutigam 1987 that airlines overprovided 'quality' and/or service in the years before deregulation.

The structural results seem to indicate a high relative cost of engine removal plus shop visit, irrespective of the era of operation or assumed discount factor; estimated values are between

0.7 and 0.9 (except for a few of the JT8D firm-specific estimates). In all but one of the regimes, the shop visit component is clearly the largest relative component of the cost vector, as much as 32 times as costly as an engine shutdown and 96 times as costly as 795 hours of operation. Thus, the maintenance problem facing an airline is an important one, with consequences to both management and society at large.

Another interesting result is the statistically significant decline in the the relative perceived cost of engine shutdown from regulation to deregulation for JT9D engines (and that the cost seems to be negative for several of the airlines when the model is estimated separately).^{11,12} An interpretation of this result is that nonfailure engine shutdowns may have become more prevalent since deregulation; that is, a larger percentage of shutdowns are caused by faulty indicator equipment rather than engine problems. Unfortunately, we have no data available to test this hypothesis. However, anecdotal evidence from mechanics suggests that the incidence of 'placarding', or citations from the FAA for minor infractions, has increased, which squares with the notion of more frequent indicator equipment failure.

Some evidence on this issue can be inferred from the results. Figure 1 shows a plot of the normalized expected value function for each engine type under each regulatory regime, with the lower curve representing operating hours with a shutdown history and the higher curve representing operating hours without a shutdown history. For JT9D engines during regulation, there appears to be little gain from a shutdown-free engine history, whereas after deregulation the gap between the relative value functions widens considerably. The reverse seems to be true for the JT8D engines, although the difference is not as great. The explanation of an increased willingness to accept placarding seems to be consistent with the JT8D results, but not with those of the JT9D, perhaps due to the more frequent deployment of JT9D engines on aircraft flying internationally (and perhaps therefore subject to more stringent inspections).

Relative operating costs increased since deregulation for JT8D engines but decreased for JT9D engines. Both changes are statistically significant. This result may be valid in and of itself, because JT8Ds are an older design than JT9Ds, or may be an artifact of noisy data that fail to distinguish the operating cost coefficient for JT8Ds from zero during the regulated years (cost specifications other than linear also resulted in nonstatistically significant estimates for θ_1). Otherwise, the results fail to explain why JT8D engines have higher relative operating costs during the unregulated era but fewer shop visits—if operating costs rose, the number of shop visit should have increased, not decreased.


The change in shop visit cost from regulated year to unregulated years is not significant for JT8D engines, but is positive (and significant) for JT9Ds. This result is consistent with changes of opposite sign in relative operating cost.

Figures 2–5 are plots of the parametric versus the nonparametric 'hazard'. Note that the shutdown hazard curve lies above the no-shutdown curve. These plots indicate one problem with this type of model. There are relatively few data points in the region where the state variable is large. Thus, the model predicts the observed outcome reasonably well for the first 10 states or so, but not in higher values of the state variable.


A major advantage to a structural rather than a reduced-form model lies in the ability to simulate the long-run behaviour of the airlines based on the assumed knowledge of the

¹¹ Statistical significance here refers to the value of the approximately normally distributed statistic used to test the null hypothesis that two parameters are equal. Significance is determined at the 10% level as in the rest of this paper.

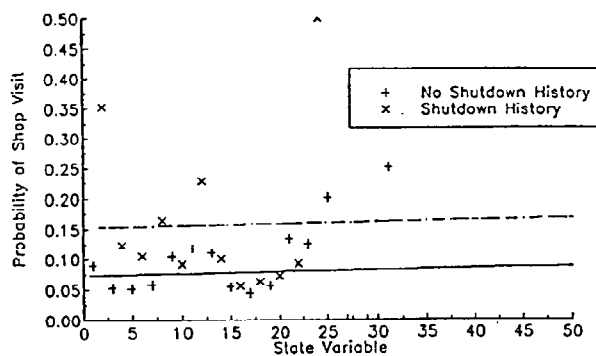
¹² Statistical significance in the cases where a negative coefficient is reported should be taken with a heavy dose of salt, because the standard errors for nearly all the firm-specific estimates are computed using the Broyden–Fletcher–Goldfarb–Shanno estimate of the Hessian matrix, rather than either the actual Hessian or the scoring matrix (as is more usual). The more usual forms were found to be singular and thus could not be inverted.



(a)



(b)

Figure 1. Expected value functions for (a) JT8D and (b) JT9D engine ($\beta = 0.9999$)Figure 2. Parametric and nonparametric hazards for JT8D engines, regulated years ($\beta = 0.9999$)

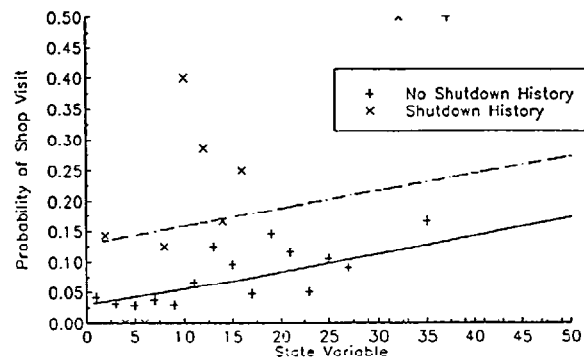


Figure 3. Parametric and nonparametric hazards for JT8D engines, unregulated years ($\beta = 0.9999$)

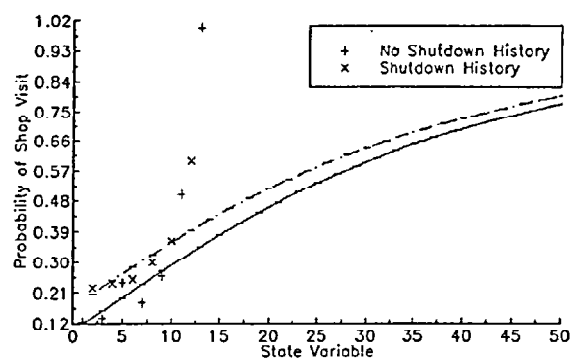


Figure 4. Parametric and nonparametric hazards for JT9D engines, regulated years ($\beta = 0.9999$)

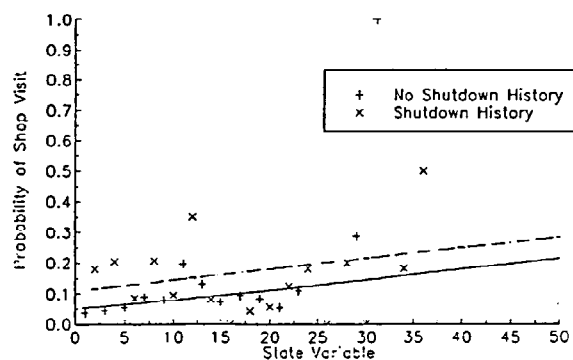


Figure 5. Parametric and nonparametric hazards for JT9D engines, unregulated years ($\beta = 0.9999$)

structural parameters. For example, we can infer the demand function for engine shop visits. Let $p_{\theta}^*(x, d)$ be the long-run stationary distribution of the process $\{d_t, x_t\}$. If we assume that for each engine m , $\{d_t^m, x_t^m\}$ is independent of $\{d_t^k, x_t^k\}$ for m not equal to k , then expected demand for shop visits D is

$$D(SVC) = MT \int_0^{\infty} p_{\theta}^*(dx, 1) = MT \sum_{x=0}^{87} p_{\theta}^*(x, 1) \quad (15)$$

where M is the number of engines and T is the number of month. By varying SVC , it is possible to plot the demand curve. Figures 6 and 7 show demand for shop visits for JT8D and JT9D engines, respectively, under regulated and unregulated regimes. Note that the shop visit cost has been normalized to the approximate average cost of a shop visit: \$150,000 and \$350,000 for JT8D and JT9D engines, respectively.¹³ We observe that the JT8D demand during regulation lies strictly above the curve since deregulation over most of the domain, and that JT9D demand before deregulation lies strictly above demand since deregulation everywhere.




Figure 6. Structural simulated demand function for JT9D engines

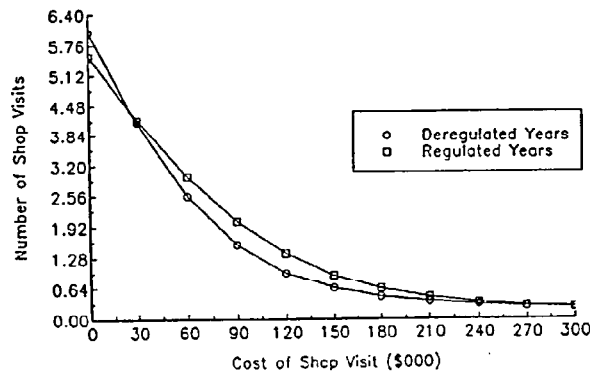


Figure 7. Structural simulated demand function for JT8D engines

¹³ According to Pratt & Whitney, Inc., personal correspondence.

In Figures 8 and 9, we can see the corresponding simulated demand functions from the reduced form estimation reported in Table VI; they appear roughly similar to the structurally derived one.

Similarly, the long-run stationary distribution can be used to compute both the expected number of hours on an engine and the overall probability of experiencing at least one shutdown. The former is simply the sum of even states times each state's associated marginal probability that is

$$\text{Expected number of hours} = 795 \sum_{j=0}^{87} \{ \text{trunc}(j/2) \{ p_{\theta}^*(j, 1) + p_{\theta}^*(j, 0) \} \} \quad (16)$$

We may be more interested in operating hours conditioned on the shop visits decision. The conditional expectations are:

$$E(\text{hours} | d = 0) = 795 \sum_{j=0}^{87} \left\{ \frac{\text{trunc}(j/2) p_{\theta}^*(j, 0)}{\sum_{i=0}^{87} p_{\theta}^*(i, 0)} \right\} \quad (17)$$




Figure 8. Reduced-form demand function for JT8D engines




Figure 9. Reduced-form simulated demand function for JT9D engines

and

$$E(\text{hours} \mid d = 1) = 795 \sum_{j=0}^{87} \left\{ \frac{\text{trunc}(j/2) p_{\theta}^*(j, 1)}{\sum_{i=0}^{87} p_{\theta}^*(i, 1)} \right\} \quad (18)$$

Note that the terms in the denominators in equations (17) and (18) are the marginal probabilities of the shop visit decision being conditioned on. The overall shutdown probability is

$$\Pr(\text{shutdown}) = \sum_{j=0}^{43} \{p_{\theta}^*(2j+1, 1) + p_{\theta}^*(2j+1, 0)\} \quad (19)$$

Again, the probability of shutdown conditioned on the shop visit decision may be of interest. The conditional shutdown probabilities are

$$\Pr(\text{shutdown} \mid d = 0) = \sum_{j=0}^{43} \frac{p_{\theta}^*(2j+1, 0)}{\sum_{i=0}^{87} p_{\theta}^*(i, 0)} \quad (20)$$

and

$$\Pr(\text{shutdown} \mid d = 1) = \sum_{j=0}^{43} \frac{p_{\theta}^*(2j+1, 1)}{\sum_{i=0}^{87} p_{\theta}^*(i, 1)} \quad (21)$$

In equations (19)–(21) it is important to recall that the probabilities are those for a history of shutdowns, i.e. the event that at least one shutdown has occurred since the last shop visit.

Table VII gives results for the conditional expectations, conditional probabilities, and overall probabilities for the engines, with β fixed at 0.999. The table confirms the nonparametric results in Kennet (1993) in showing that expected hours of operation at shop visit increased for both engine types as a result of deregulation. However, the means conditional on $d = 1$ do not closely match the values computed nonparametrically, a fact perhaps attributable to error in the discretization process.

The stationary distribution results for probability of shutdown history conditional on shop visit also do not closely match nonparametric estimates. The latter values are 0.0191, 0.1290, 0.0732, and 0.1036, for JT8D regulated, JT8D unregulated, JT9D regulated, and JT9D unregulated, respectively. However, the parametric estimates are heuristically consistent with the logit results in Kennet (1993), which show either a negative or statistically insignificant

Table VII. Long-run stationary distribution results

Item	JT8D regulated	JT8D unregulated	JT9D regulated	JT9D unregulated
$E(\text{hours} \mid d = 0)$	2808.4	3678.3	1647.5	3385.9
$E(\text{hours} \mid d = 1)$	3028.1	4666.1	2031.6	4257.4
$\Pr(\text{shutdown} \mid d = 0)$	0.1050	0.0216	0.1739	0.1656
$\Pr(\text{shutdown} \mid d = 1)$	0.2106	0.0726	0.2894	0.3179
$\Pr(\text{shutdown})$ (unconditional)	0.1138	0.0243	0.1940	0.1782

relationship between deregulation and the probability of shutdown after correcting for other factors.¹⁴

In any case, the predicted outcomes are long-run stationary results, which may or may not be achieved in the relatively short-run time frame of the dataset.

5. CONCLUSION

Earlier work established that a dataset containing complete Pratt & Whitney aircraft engine histories exhibits a clear distinction between maintenance behaviour before and after airline deregulation. In an effort to learn what caused the policy change we develop and estimate a plausible structural model of aircraft engine maintenance. We estimate structural parameters separately for the regulated and unregulated eras for each engine type. A likelihood ratio test confirms the validity of separating the sample, which verifies the earlier nonparametric and parametric reduced-form results in Kennet (1993).

An interesting result that emerges is that the data seem to support the notion of dynamically optimizing maintenance scheduling during the era since deregulation, while rejecting the notion for the earlier regulated period in the sense that we cannot distinguish alternative specifications in which there is little discounting and in which the future is completely discounted. Additionally, the data exhibit firm heterogeneity in the deregulated era, while not in the regulated era, suggesting that firms have become more independent in constructing maintenance strategies.

The structural model results offer a possible explanation for the observed relationship. The level of maintenance before deregulation was superoptimal in that firms did not seem to minimize discounted costs. Deregulation, in this scenario, forced firms to become more cost efficient, driving maintenance effort to a more efficient level. Supporting this claim is the result that deregulation seems to have had no effect on the likelihood of an engine shutdown.

The data also seem to indicate that the perceived cost to firms of engine shutdown is lower since deregulation. For at least one of the engine types, the results are consistent with claims by workers of corporate laxity toward minor airworthiness violations.

ACKNOWLEDGEMENTS

Thanks are due to Arthur Wegner, Edward Cowles, and Robert Weiss of Pratt & Whitney, Inc., for permission to use their data and to Charlie Walters, Roy Fife, and Johnny Townsley of Federal Express, Inc., for telling the author what the data meant. Thanks also to Mita Das, Zvi Eckstein, co-editor Nick Kiefer, an anonymous referee, Marshall Reinsdorf, Leigh Riddick, Jim Spletzer, İnsan Tunali, Michael Zwecher, and especially John Rust for their help on the earlier versions of this study, and to Julie Kenkel for research assistance. The author places all blame for remaining errors squarely on the shoulders of Microsoft, Inc., for not having invented a Thought Checker.

¹⁴We should point out that a direct comparison is not valid because the logit estimation was performed unconditionally with respect to shop visit and refers only to the probability of a shutdown occurring in a particular month (rather than to the probability of a history of shutdowns).

REFERENCES

- Billingsley, P. (1979), *Probability and Measure*, John Wiley, New York, pp. 309–310; p. 320.
- Braeutigam, R. (1987), *Transportation Deregulation and Safety*, Conference Proceedings, Transportation Center, Northwestern University, 23–25 June.
- Das, S. (1992) 'A micro-econometric model of capital utilization and retirement: the case of the U.S. cement industry', *Review of Economic Studies*.
- Golbe, D. (1986), Safety and profits in the airline industry, *The Journal of Industrial Economics*, XXXIV, No. 3, March.
- Kamien, M. I. and D. Vincent, (1990) 'The effects of price regulation and consumer ignorance on quality of service', Mimeo, Department of Managerial Economics and Decision Science, Kellogg School of Business, Northwestern University.
- Kennet, D. M. (1993), Did deregulation affect aircraft engine maintenance? An empirical policy analysis, *The RAND Journal of Economics*, 24, No. 4, Winter.
- Panzar, J. and I. Savage (1987), *Transportation Deregulation and Safety*, Conference Proceedings, Transportation Center, Northwestern University, 23–25 June.
- Press, W. H., B. P. Flannery, S. A. Teukolsky and W. T. Vetterling (1986), *Numerical Recipes: The Art of Scientific Computing*, Cambridge, Cambridge University Press, Chapter 10.
- Rose, N. L. (1981), 'Financial influences on airline safety', in L. Moses and I. Savage (eds), *Safety Performance Under Deregulation*, Oxford University Press, Oxford.
- Rose, N. (1990) 'Profitability and product quality: economic determinants of airline safety performance,' *Journal of Political Economy*, 98, No. 5, October.
- Rose, N. (1992), 'Fear of flying? Economic analyses of airline safety', *Journal of Economic Perspectives*, 6, No. 2, Spring.
- Rust, J. (1987), 'Optimal replacement of GMC bus engines: an empirical model of Harold Zurcher', *Econometrica*, 55, No. 5, September.
- Rust, J. (1988), Maximum likelihood estimation of discrete control pocesses', *SIAM Journal of Control and Optimization*, 26-5, 1006–1024.
- Schwartz, A. N., J. A. Sheler and C. R. Cooper (1971), Dynamic programming approach to the optimization of naval aircraft rework and replacement policies', *Naval Research Logistics Quarterly*, 18, 395–414.
- Sherif, Y. S. and M. L. Smith (1981), 'Optimal maintenance models for systems subject to failure—a review', *Naval Research Logistics Quarterly*, 32, 47–74.
- Sherwin, D. J. (1979), 'Inspection intervals for condition-maintained items which fail in an obvious manner', *IEEE Transactions on Reliability*, R-28, 1, 85–89.
- Stiglitz, J. and R. Arnott (1987), *Transportation Deregulation and Safety*, Conference Proceeding, Transportation Center, Northwestern University, 23–25 June.
- Sturm, R. (1990), 'A structural economic model of operating cycle management in European nuclear power plants', Mimeo, Stanford University.

This document was created with Win2PDF available at <http://www.daneprairie.com>.
The unregistered version of Win2PDF is for evaluation or non-commercial use only.