On the Optimal Lifetime of Nuclear Power Plants

Geoffrey Rothwell
Department of Economics, Stanford University, Stanford, CA 94305 (rothwell@leland.stanford.edu)

John Rust
Department of Economics, Yale University, New Haven, CT 06520-8264 (jrust@econ.yale.edu)

We present an empirical model of optimal operation of nuclear power plants. The optimal lifetime is the solution to an optimal stopping problem: The plant is closed when the expected discounted losses from continued operation exceed the discounted costs of decommissioning. We forecast the evolution of the nuclear power industry under the current regime of 40-year operating licenses and for a policy allowing 20-year license extensions. We conclude that the extension would double the expected discounted value of U.S. nuclear power plants and double the undiscounted electrical power output of the U.S. nuclear industry over its remaining lifetime.

KEY WORDS: Atomic Energy Act; Dynamic programming; Electricity generation; Life extension; Maximum likelihood estimation; Nuclear Regulatory Commission.

Much attention has been paid to the decline of the nuclear power industry in the United States following the Three Mile Island (TMI) accident in March 1979. Due to large increases in construction costs, lead times, and operating expenses induced by stricter regulation of the nuclear power industry, there have been no new orders for nuclear power plants (NPP's) since 1978. Orders for more than 100 NPP's were cancelled—some involving plants that were nearly complete—resulting in losses of tens of billions of dollars in planning and construction costs (Energy Information Administration 1983). In a previous article (Rust and Rothwell 1995), we estimated that more than 90% of the discounted profits to continued operation of existing nuclear power plants were eliminated in the stricter regulatory regime after the TMI accident.

Despite the reduction in profits, more than 100 NPP's continue to operate in the United States, constituting a combined generating capacity of 100 gigawatts or about 15% of the country's total electrical capacity. In 1992 these nuclear power plants produced 620 billion kilowatt-hours, or 22% of the nation's electricity supply (Institute for Nuclear Power Operations 1993). There will be a significant loss of generating capacity over the next three decades, however, from retirements of aging NPP's. Some studies (e.g., Forest, Deutsch, and Schenler 1988) have estimated that extending the life of existing nuclear power plants by 20 years could result in large savings—about \$450 billion in current dollars.

Under the Atomic Energy Act, the Nuclear Regulatory Commission (NRC) issues operating licenses for a maximum term of 40 years. In 1991 the Commission developed a preliminary set of procedures that would allow NPP's to apply for an extension of their operating licenses by an additional 20 years. Although the Commission is currently solidifying its procedures for extending operating licenses, the combination of increased operating costs and a variety of age-related problems have led to the presumption that license extension would not be economic. For example, Hewlett (1991) concluded that, unless there is a reduction in both the level and the rate of growth in operating and maintenance costs experienced by NPP's in the mid-1980s, the discounted costs of a 20-year NPP life-extension program "are roughly equal to or greater than the cost of constructing other types of power plants" (p. 271). It is not even clear that utilities will find it profitable to continue operating existing NPP's for the full 40-year duration of their initial license. Indeed, since 1988 six NPP's have been closed more than 10 years before their license expiration dates. These closures could foretell of many more: "Some analysts suggest that as many as 25 plants, not necessarily older ones, may be found uneconomic during the next several years" (Office of Technology Assessment 1993, p. 21).

This article presents an empirical model of an operator's decision whether to operate or close an NPP. Our model extends a dynamic programming (DP) model of optimal NPP operation developed by Rust and Rothwell (1995) by allowing for the occurrence of "major problem spells." Examples of major problem spells include extended shutdowns to repair or replace major reactor components, such as steam generators or damaged reactor vessels, as well as extended NRC-mandated outages to correct management or safety problems. The DP model predicts that under ordinary operating conditions it is unlikely that a plant will be closed, but the probability of closure for decommissioning increases substantially during a major problem spell.

We use the estimated DP model to forecast nuclear power generation under two policy scenarios, (1) a fixed 40-year license span with no possibility of extension and (2) a "costless" extension in operating licenses to 60 years. The DP model shows that the length of the operating license critically affects the economics of the closure decision. Under a 40-year license, the remaining horizon is too short to recoup the costs incurred for a major retrofit or major problem spell that occurs beyond the 24th year of the operating license, so whenever a major problem occurs beyond the 24th year, the optimal decision is to close the NPP. Interestingly, the DP model also predicts that it is optimal to close the plant if a major problem occurs during the very first year of a plant's life. The optimal policy is different under a 60-year license span. In this case the remaining horizon is long enough to recoup the costs of a major problem spell in all but the last 12 years of the operating license.

Section 1 provides a brief overview of nuclear power generation and regulation and discusses some of the age-related deterioration problems experienced by NPP's and the uncertainties related to applying for 20-year license extensions. Some background on these issues is necessary to understand our specification of the DP model and to motivate the practical issues underlying the policy simulations in Section 4. Section 2 presents our DP model of optimal operation of an NPP. In the DP model the operator must decide each month whether to operate, refuel, or close the NPP. The operating decisions depend on the signals the operator receives about the NPP's operating status, some of which are recorded in our dataset and some of which are unobserved by the econometrician. Our DP model accommodates both types of signals. The dataset used to estimate the DP model is based on information available in the NRC's (1989–1994) NUREG-0020, commonly known as the Graybook. Although this dataset contains detailed monthly observations of NPP operating decisions, we have no data on NPP operating costs at monthly time intervals. Nevertheless, our "revealed preference" econometric methodology enables us to infer the ratio of the change in expected discounted profits of NPP operations resulting from a 20-year extension in operating licenses.

In Section 3 we use the Graybook data to estimate the unknown parameters of the electric utility's profit function, the failure processes that lead to unplanned forced outages, and the parameters governing the duration of refueling outages. Section 4 presents simulation results from our estimated DP model, including predictions of industry output under the two different licensing scenarios summarized previously. Section 5 presents concluding remarks and directions for future research.

1. NUCLEAR POWER TECHNOLOGY, PLANT AGING, AND PROSPECTS FOR LICENSE EXTENSION

There are many types of NPP's, but in the United States nearly all commercial NPP's are light water reactors (LWR's), which use ordinary water as coolant and moderator. Nearly all commercial LWR's in the United States are one of two types, pressurized water reactors (PWR's) or boiling water reactors (BWR's). Of the 111 licensed U.S. NPP's operating in 1993, 76 were PWR's and 35 were BWR's. LWR's generate power via nuclear fission using slightly enriched uranium in numerous zirconium-cladded fuel rods that are inserted into the reactor core in bundled fuel assemblies. In BWR's, the coolant flow that is pumped into the reactor vessel is allowed to vaporize after coming in contact with the heated fuel assemblies, and the resulting (slightly radioactive) steam is piped directly to steam turbines driving electrical generators. PWR's have a two-loop system with an isolated inner loop of pressurized, superheated coolant that runs through the reactor vessel. The superheated coolant from the inner loop transfers the reactor's heat energy to an outer cooling loop via *steam generators,* which deliver the steam that runs the turbines attached to the electrical generators. The advantage of the two-loop system in a PWR is that the steam running through the turbines is not radioactive as in the case of BWR's. The disadvantage is that the thermal efficiency of a PWR is generally lower than that of a BWR (32% vs. 34%).

The high energy released by fission has deleterious effects on the structure of the fuel rods. Some fission products appear as gases that eventually create pressure within the fuel rods. As a result, a fuel rod can swell, crack, and become physically distorted to such an extent that it is no longer usable. The loss in fuel reactivity due to gradual depletion of radioactive uranium and buildup of fission products, combined with the effect of radiation-induced fuel swelling and distortion, are limiting factors determining how long an LWR can run between refuelings. The maximum safe duration between refuelings is a function of the initial level of enrichment of the uranium, the design of the fuel rods, and the fuel-management strategy adopted by the operator.

One of the difficult problems confronting plant operators is to determine the optimal length of operating (or refueling) cycles. There is a primary trade-off between (a) the potential improvement in capacity factor associated with longer operating cycles and (b) the potential increased risk of unplanned mid-cycle outages due to fuel and other failures. We consider how this trade-off changes as NPP's age. NPP's experience two sorts of aging problems, (1) short-term or "within-cycle" aging problems such as fuel-rod failures that increase with the duration of operating spells and (2) long-term or "between-cycle" aging problems such as radiation-induced deterioration of the reactor vessel or corrosion problems in steam generators of PWR's. Refueling outages partially regenerate the within-cycle deterioration associated with the burnup of nuclear fuel.

It is more difficult to quantify the impact of long-term aging problems because plant-specific learning-by-doing effects (Lester and McCabe 1993) and general technological improvements in fuel reliability, instrumentation, and other aspects of nuclear power technology have more than offset the between-cycle depreciation of NPP components leading to steady improvement in overall NPP performance as reactor age increases. In particular, the rate and duration of unplanned outages decrease monotonically with NPP age (e.g., see Rothwell and Rust 1995; Rothwell 1996). Part of the difficulty in identifying the independent effect of age-related deterioration is that reactor age is highly correlated with calendar time because two-thirds of U.S. reactors came online during the late 1960s and 1970s. It is possible that we will observe an acceleration in age-related degradation in NPP's toward the end of their operating life, but unfortunately there are few observations on U.S. NPP's that are more than 20 years old. Another "problem" is that safety procedures are designed to ensure that failures are rare events, making it difficult to estimate their hazard rates from limited observations. The policy forecasts made in this article assume that regular maintenance including large-scale capital upgrades such as replacement of corroding steam generators in PWR's succeed in preventing any sudden deterioration in NPP components toward the end of the plant's 40-year operating licenses. Indeed, we assume that with regular preventive maintenance a plant can operate safely indefinitely. In addition we assume that the historical rate of technological improvement in fuel reliability and other NPP operating components will continue for the foreseeable future. Under these assumptions, it is not necessary to separately identify the independent effects of long-term aging and technological improvements: We only need to identify the net effect in our forecasts of NPP operating performance over their remaining lifetimes.

There is a large engineering literature on problems and strategies for counteracting age and utilization-related deterioration in NPP's (e.g., see Shah and MacDonald 1993). Through direct measurements of individual reactor components, this literature has identified key aging mechanisms and developed appropriate age-management strategies for dealing with them. Shah and MacDonald's (1993) ranking of the primary degradation sites of the major LWR components identified damage to the reactor vessel from radiation embrittlement and boric-acid corrosion as one of the most important age-related safety hazards. Embrittlement of the reactor vessel creates a significant potential safety hazard due to a phenomenon known as *pressurized thermal shock:* Cold water entering the coolant stream could cause a sudden lowering of the temperature in the reactor vessel causing crack initiation, propagation, or fracture. Failures in the main reactor components are so expensive to repair that discovery of these problems can precipitate the closing of the plant. Although some studies have claimed that embrittlement problems can be reversed by *thermal annealing,* there has been little practical experience with this procedure and there is substantial scientific uncertainty about the rate of reembrittlement after annealing (Shah and MacDonald 1993, p. 64). Thus, the main strategy for dealing with these problems is through preventive maintenance and conservative operating practices, including "low leakage" fuel-management strategies. Embrittlement of reactor vessels is one of the main public-safety uncertainties affecting the NRC's policy of operating-license extension. It is an example of one of the "major problem" factors that could precipitate a closure of an NPP: In fact, embrittlement was a significant factor behind the 1991 closure of the Yankee Rowe reactor.

Reactor embrittlement is just one of many factors that could lead to closure of an NPP, however: The decision to retire a plant depends critically on the overall reliability and cost-efficiency of the plant, the length of the operating license (which determines the horizon over which major maintenance investments can be recouped), and the costs involved in plant decommissioning. Under the Atomic Energy Act, the NRC is allowed to issue operating licenses for a maximum term of 40 years. Once an NPP is retired (either at the expiration of its operating license or due to early retirement for economic or safety reasons), NRC regulations require that decommissioning be performed to protect the public and the environment from exposure to radioactivity. To date in the United States, only small NPP's have been completely decommissioned. Current estimates of decommissioning costs for a typical 1,000-megawatt NPP are in the range of several hundred million dollars (see Pasqualetti and Rothwell 1991; Office of Technology Assessment 1993). In this article we assume that the least-cost decommissioning alternative is chosen.

In 1991 the NRC began to draft a set of procedures that would allow operators the opportunity to apply to extend their operating licenses by as much as 20 years. The NRC license-extension policy imposes a severe burden of proof on the applicant to show that the plant can be safely operated beyond its original license term. This is in contrast to regulatory regimes with unlimited-term licenses in which the burden is on the regulator to prove that a plant should have its license revoked because it can no longer be operated safely. Because the criteria governing NRC license-extension decisions are not yet finalized, there is considerable uncertainty about the conditions under which license extensions will be granted. The NRC has estimated that a typical license-renewal application would require "approximately 200 person-years of utility effort (supplemented by unquantified consultant support) and span 3 to 5 calendar years at a cost of about $30 million" (Office of Technology Assessment 1993).

Because of the complexities and uncertainties surrounding the slowly evolving NRC license-renewal procedures, many NPP operators believe that license renewal is not a realistic possibility. As of 1996, no operator has applied for a license extension from the NRC. In late 1992, the owner of the Monticello NPP indefinitely postponed its plans to submit a license-renewal application, citing concern about the interpretation of the NRC's rule and noting that the number of systems to be reviewed had grown from the original 74 to 104 with "no indication of where it might go from here" (p. 59, Office of Technology Assessment 1993). As we show in Section 4, if NPP operators believe that there is no realistic possibility of license extension, there is a high probability that NPP operators will actually close their plants well before the end of their 40-year license span.

In view of the large costs involved in "premature closure" of U.S. NPP's, the Office of Technology Assessment (1993) suggested that it would be wise to resolve the uncertainties surrounding the NRC's current license renewal procedures: "If ongoing aging management programs are adequate during the original license term, it may be possible to considerably simplify the license renewal rule without affecting safety. ... For this reason, it may be better to view aging management as a more continuous process than established in the license renewal rule" (p. 18). A logical extension of this line of reasoning would be for the NRC to adopt a "costless" license-extension rule in which the term of each licensed NPP would automatically be extended to 60 years unless the NRC regulators determined that the plant could not continue to be operated safely. We argue that the appropriate policy regarding operating-license extensions can be usefully analyzed within the context of a DP model of NPP operations. We introduce such a model in Section 2. In Section 4 we use the DP model to show how costless license extension would lead to a substantial increase in the expected discounted profits of U.S. NPP's and thereby extend the lifetime of the nuclear power industry. We emphasize, however, that, although our DP model allows us to quantify the benefits of a policy of a costless 20-year extension in operating licenses, it does not allow us to quantify all the potential costs involved. The most significant cost is the potential threat to public safety posed by 100 aging NPP's whose reactor vessels could become dangerously embrittled during an additional 20 years of regular operations. As we noted previously, it is very difficult to forecast the rate of embrittlement, and it is equally difficult to quantify the risk and potential losses stemming from a major nuclear accident such as the rupture of a reactor vessel due to the combined effect of embrittlement and pressurized thermal shock.

2. A DYNAMIC PROGRAMMING MODEL OF NPP OPERATIONS

In this section we present a DP model designed to capture the main features of NPP operations discussed previously. The DP model is based on the maintained hypothesis that NPP operators control their plants to maximize expected discounted profits from electricity generation subject to technological and regulatory constraints. Although most NPP's are owned by utilities who are typically modeled as monopolists subject to rate-of-return regulation, empirical evidence presented by Rust and Rothwell (1995) suggests that profit maximization provides a reasonable description of utilities' objectives in the post-TMI accident regulatory environment. This is because of the sharp rise in the probability of operating-cost disallowances in public utility commission (PUC) ratemaking decisions, and the increasing prevalence of other incentive-based regulations has encouraged profit-maximizing behavior on the part of utilities—see Che and Rothwell (1995). In addition, falling costs of fossil fuels during the 1980s put strong pressure on utilities to operate their NPP's in a cost-minimizing manner, subject to the more stringent safety regulation imposed by the NRC that we do account for in our model.

Furthermore, we presume that utilities are able to freely buy and sell electricity over power grids, so NPP's are not subject to additional generation constraints from low local demand for electricity or lack of reliable sources of replacement power. This implies that the current price of electricity is the relevant shadow price governing plant operating decisions, enabling us to treat each NPP as a separate profit center and abstract from having to model idiosyncratic fluctuations in local demand and the operating status of other generating units owned by the utility. It also justifies our use of observations on NPP *availability factors* (the fraction of the month the NPP is operating) rather than *capacity factors* (the ratio of energy actually generated during a month to the energy that would have been generated if the NPP was run continuously at maximum dependable capacity). The availability factor is greater than the capacity factor when a plant is operating at less than full utilization due to "load following"—that is, reducing power output to meet local demand constraints. Because most NPP's are baseloaded and attached to power grids, they are not frequently subject to local demand constraints, so availability and capacity factors typically coincide. As a result, we do not expect that our results would change significantly if we were to reestimate the DP model using observed capacity factors. [See Rothwell (1990) for a more detailed discussion of the difference between availability or service factors.] Our model does account for seasonal fluctuations in power demand and its implications for operating and shutdown decisions of NPP's, however.

Moreover, unlike Rust and Rothwell (1995), we estimate the DP model for a sample from January 1989 to December 1994. NPP operating and maintenance (O&M) costs rose rapidly (at 11% per year from 1974 to 1984 and 5% per year from 1985 to 1989), peaking in 1989 when they exceeded the O&M costs for coal plants for the first time (Office of Technology Assessment 1993, p. 24; Energy Information Administration 1995, p. 29). Since 1989, however, real O&M costs have stabilized. Therefore, we restrict our sample to 1989 through 1994 when the growth in O&M costs stabilized and industry capacity factors had returned to the high levels observed before the TMI accident.

In this article we extend the DP model of Rust and Rothwell (1995) to a three-spell model. We define a *major problem spell* as any continuous shutdown that lasts longer than nine months. Major problem spells are infrequent events (there are 23 such spells in our dataset over the period 1989 to 1994), and they occur for a variety of reasons, including overhaul or replacement of major reactor components (such as steam generators in a PWR) and administrative reasons. Both the incidence and duration of many of these major problem spells are involuntary and beyond the direct control of the operator. Because major problem spells are rare events, we do not attempt to distinguish between administrative shutdowns and shutdowns required to undertake major repairs or equipment backfits. We assume that major problem spells are exogenous stochastic events; that is, the operator lacks control over both their incidence and duration. Due to the high costs of major problem spells—both the direct capital additions and maintenance costs and the opportunity costs of lost power generation—we will show that once a plant enters a major problem spell the chances that it will be closed for decommissioning increase substantially.

We follow Rust and Rothwell (1995) and formulate the DP problem in discrete time. NPP operating decisions are assumed to be made at the start of each month. In reality, the operator must control the NPP in continuous time. Our model abstracts from the details of the minute-by-minute decisions made by the operator. Although day-to-day management of NPP's involves complex trade-offs, given the high opportunity costs of NPP downtime, these complexities are secondary to the larger long-run issues of operating NPP's. Our DP model is designed to capture the most important longer-run trade-offs that NPP operators face—that is, the duration of operating spells, the timings of

Table 1. State and Control Variables of DP Model

State variables: $x_t = (r_t, f_t, d_t)$, where
- r_t = type of spell in *previous* month
  - $r_t = 1$ if the previous month was part of a major problem spell
  - $r_t = 2$ if the previous month was part of a refueling spell
  - $r_t = 3$ if the previous month was part of an operating spell
- f_t = NPP signal in *current* month
  - $f_t = 1$ no signals that require initiation of a forced outage are received during the month
  - $f_t = 2$ operator receives signals requiring one or more forced outages
  - $f_t = 3$ if $r_t = 3$ operator observes "enter major problem spell" signal
    if $r_t = 2$ operator observes "continue refueling spell" signal
    if $r_t = 1$ operator observes "continue major problem spell" signal
- d_t = duration of spell in *previous* month
  - $r_t = 1$ d_t is duration of major problem spell
  - $r_t = 2$ d_t is duration of refueling spell
  - $r_t = 3$ d_t is duration of operating spell

Control variable a_t:
- If $r_t = 3$ and $f_t < 3$, the NPP is in an operating spell and the operator has not received a major problem signal and the choice set is $A_t(x_t) = \{1, \ldots, 8\}$, given by
  - $a_t = 1$ permanently close the NPP
  - $a_t = 2$ refuel the NPP
  - $a_t = 3$ shut down the NPP (i.e., run the NPP at 0%)
  - $a_t = 4$ run the NPP between 1% and 25% availability
  - $a_t = 5$ run the NPP between 26% and 50% availability
  - $a_t = 6$ run the NPP between 51% and 75% availability
  - $a_t = 7$ run the NPP between 76% and 99% availability
  - $a_t = 8$ run the NPP at 100% of its potential output
- If $r_t = 3$ and $f_t = 3$, the NPP is in an operating spell and the operator receives a major problem signal and the choice set is $A_t(x_t) = \{1, 2, 3\}$.
- If $r_t = 1$ and $f_t = 3$, the NPP is in a major problem spell and the operator receives a signal that the major problem spell will continue for one more month and the choice set is $A_t(x_t) = \{1, 3\}$.
- If $r_t = 2$ and $f_t = 3$, the NPP is in a refueling spell and the operator receives a signal that the refueling spell will continue for one more month and the choice set is $A_t(x_t) = \{1, 2\}$.
- If $r_t < 3$ and $f_t < 3$, the NPP is in a refueling spell or a major problem spell and the operator receives a signal that the refueling spell or major problem spell has ended and the choice set is $A_t(x_t) = \{1, 3, 4, 5, 6, 7, 8\}$.

planned outages for refuelings and preventive maintenance, and whether a plant should be closed for decommissioning.

The DP model consists of a vector of *state variables* s_t, a *control variable* a_t, a *profit function* $\pi(a, s)$, a *discount factor* β, and a *transition density* $\lambda(s'|s, a)$, representing the stochastic law of motion for the state of the plant. A key advantage of the DP framework is that it allows us to determine optimal operating strategies that account for uncertain events like the occurrence of forced outages and major problems. The operator clearly observes more signals about the NPP's current operating status than are available in our dataset or that are even feasible to record. Therefore, we assume that the state variable s_t can be partitioned into two components, $s_t = (x_t, \varepsilon_t)$, where x_t is an *observed state vector* and ε_t is an *unobserved state vector*. The operator observes both components, but we observe only x_t. The NPP operator weighs the consequences of various operating decisions given the full set of signals and takes the best action. We assume that the result of this decision process can be summarized by a vector of current net benefits (or costs, if negative) to each operating decision. Thus, we will interpret ε_t as a vector with the same number of elements as possible values of the control variable a_t. Because the full set of information available to the NPP operator is unobserved, we treat ε_t as a latent random vector with an extreme value distribution.

We follow the general framework of Rust (1987, 1988, 1995) and assume that the operator's current period profit from taking action a for the plant in current state (x, ε) is given by the function $\pi(a, x, \varepsilon)$ with the additively separable representation

$$\pi(a, x, \varepsilon) = \mu(a, x, \phi) + \varepsilon(a), \tag{2.1}$$

where ϕ is a vector of unknown profit-function parameters to be estimated. We assume that the vector of state variables (x, ε) evolves according to a controlled Markov process with transition density $\lambda(x_{t+1}, \varepsilon_{t+1}|x_t, \varepsilon_t, a_t)$ and that the NPP operator chooses an optimal operating strategy $a_t = \alpha_t(x_t, \varepsilon_t)$ that maximizes the NPP's expected net present value $V_0(x, \varepsilon)$ given by

$$V_0(x, \varepsilon) = \max_{(\alpha_0, \ldots, \alpha_T)} E\left\{\sum_{t=0}^{T} \beta^t \pi(a_t, x_t, \varepsilon_t)|x_0 = x, \varepsilon_0 = \varepsilon\right\}. \tag{2.2}$$

In many DP problems, the horizon T is not well defined. For an NPP, however, the horizon, T is determined by the NRC's 40-year operating license. Therefore, we will initially assume a 40-year life, which corresponds to $T = 480$ in our monthly DP model. In Section 4, however, we consider the empirical case for other expectational hypotheses, and in Section 5 we examine the predictions of the DP model under the hypothesis that operating licenses can be costlessly extended from 40 to 60 years. The state and control variables used in the DP model are defined in Table 1.

The timing of plant signals and operating decisions is as follows: At the start of period t, the NPP operator knows the state r_t of the NPP in the previous month—that is, whether it was in a major problem spell, a refueling spell, or an operating spell. The operator also knows the duration d_t of this spell. At the beginning of the month the operator receives a signal (f_t, ε_t) summarizing the NPP's operating condition for the coming month. Conditional on this signal and the plant's state in the previous month, the operator chooses the action a_t that has the highest expected net present value of operating profits. Given a_t and (x_t, ε_t), the spell type of the current month is determined. The NPP operator updates r_{t+1} and d_{t+1} (according to rules that will be detailed shortly), new values of $(f_{t+1}, \varepsilon_{t+1})$ are realized, and the NPP operator makes the next decision in period $t+1$. Our assumption that the NPP operator observes a signal at the start of the month summarizing the NPP's status for the rest of the month is an idealization designed so that our discrete-time model could mimic the control process that occurs in continuous time. Given our interpretation of our DP model as an approximation of the continuous-time control process, we do not regard our assumptions about the timing of signals and operating decisions as reflecting "clairvoyance" by the operator. Instead, our model abstracts from the exact timing of forced outages within a month to focus attention on the more important "big picture" issues such as timing of refuelings and plant closure decisions for which a monthly interval is an appropriate level of time aggregation. We believe that the errors from our monthly approximation to the continuous-time control process are negligible in comparison to other specification errors in our model (such as the assumption that $\{\varepsilon_t\}$ is iid).

Next we specify the functional forms for the profit function $\mu(a, x, \phi)$ and the transition density $p(x'|x, a, \psi)$. The laws of motion for the state variables r_t and d_t do not require estimation:

$$r_{t+1} = \begin{cases} 1 & \text{if } f_t = 3 \quad \text{and} \quad (r_t = 3 \quad \text{or} \quad r_t = 1) \\ 2 & \text{if } (a_t = 2 \quad \text{and} \quad f_t < 3 \quad \text{and} \quad r_t = 3) \\ & \quad \text{or} \quad (r_t = 2 \quad \text{and} \quad f_t = 3) \\ 3 & \text{if } a_t > 2 \quad \text{and} \quad f_t < 3, \end{cases} \tag{2.3}$$

and

$$d_{t+1} = \begin{cases} d_t + 1\{a_t \neq 3 \quad \text{or} \quad r_t \neq 3\} & \text{if } r_{t+1} = r_t \\ 1 & \text{otherwise.} \end{cases} \tag{2.4}$$

The law of motion for duration corresponds to "partial regeneration" of the NPP following each refueling or major problem spell. In the model that we shall present, rates of forced outages and operating costs increase with the duration of the operating spell. Rates of forced outages are high immediately after a refueling, however. These initial problems are resolved rapidly, causing forced-outage rates to decline until about the 12th month of the operating spell, after which they begin to increase. In addition, Rust and Rothwell (1995) presented evidence that forced outages that occur later in the operating cycle are more "serious" in the sense that their mean durations are longer. During a refueling outage, maintenance is performed that partially regenerates the NPP. The regeneration is only "partial" because operating costs and rates of forced outages are also a function of the age of the NPP. Because of technological progress and learning-by-doing, however, the net effect of aging is estimated to be negative; that is, rates of forced outages decrease with age. Mid-cycle preventive maintenance outages also help to regenerate the NPP. This is reflected in the formula for duration, Equation (2.4): $d_t + 1$ is incremented for each period the NPP continues in the current spell ($r_{t+1} = r_t$) except when there is a maintenance shutdown during an operating spell ($a_t = 3$ and $r_t = 3$). In this case, the partial regenerative effects of the shutdown are proxied by setting $d_{t+1} = d_t$ instead of $d_{t+1} = d_t + 1$.

Plant closure is assumed to be an absorbing state: Once the operator chooses action $a_t = 1$, there are no future operating decisions to be made. Although decommissioning a plant takes time, our model will simply estimate a parameter representing the expected discounted costs incurred over the duration of the decommissioning process as a one-time charge. If the NPP has not been closed before the end of its operating license at $T = 480$, then we assume that the operator is forced to close in the final period; that is, $A_{480}(x) = \{1\}$. This assumption is relaxed in Sections 3 and 4 when we consider alternative licensing rules.

The law of motion for the NPP status variable f_t is probabilistic. Its probability distribution is derived from five conditional probabilities:

1. p_{of}—probability of one or more forced outages occurring during an operating spell
2. p_{rf}—probability of one or more forced outages occurring in the first month following a refueling outage
3. p_{om}—probability of entering a major problem spell from an operating spell
4. p_{mo}—probability of resuming operation from a major problem spell
5. p_{ro}—probability of resuming operation from a refueling outage

Each of these conditional probabilities depends on the NPP age at t and the observed state and control variables (x_t, a_t). They are estimated as binary logit probabilities given by

$$p_i(x_t, a_t, t) = \frac{\exp\{g(x_t, a_t, t, \psi_i)\}}{1 + \exp\{g(x_t, a_t, t, \psi_i)\}}, \qquad i = of,\ rf,\ om,\ mo,\ ro, \tag{2.5}$$

where g is a flexible functional form used to estimate these probabilities (typically a linear-in-parameters specification) and $\psi = (\psi_{of}, \psi_{rf}, \psi_{om}, \psi_{mo}, \psi_{ro})$ is a vector of unknown parameters to be estimated. Given these probabilities, we can define the law of motion for f_t. There are three cases to consider. If the NPP is in an operating spell (i.e., if $f_t < 3$ and $a_t > 2$), then f_{t+1} is given by

$$f_{t+1} = \begin{cases} 1 & \text{with probability } (1 - p_{om})(1 - p_{of}) \\ 2 & \text{with probability } (1 - p_{om})p_{of} \\ 3 & \text{with probability } p_{om}. \end{cases} \tag{2.6}$$

If the NPP is currently in a major problem spell or has just entered a major problem spell (i.e., if $r_t = 1$ or $r_t = 3$ and $f_t = 3$), then f_{t+1} is given by

$$f_{t+1} = \begin{cases} 1 & \text{with probability } p_{mo}(1 - p_{of}) \\ 2 & \text{with probability } p_{mo}p_{of} \\ 3 & \text{with probability } (1 - p_{mo}). \end{cases} \tag{2.7}$$

If the operator initiates a refueling outage (i.e., if $a_t = 2$) or if the current month is a continuation of a refueling spell (denoted by $r_t = 2$ and $f_t = 3$, as will be explained), there is a similar law of motion for f_{t+1} as in Equation (2.7) but with p_{ro} replacing p_{mo} and p_{rf} replacing p_{of}.

The NPP's profit function $\pi(a_t, x_t, \varepsilon_t)$ was specified to have the additive-separable decomposition in Equation (2.1). We now turn to the specification of the component $\mu(a_t, x_t, \phi)$ of π, which depends on the observed state of the NPP x_t, the operator's decision a_t, and a vector ϕ of unknown parameters to be estimated. Let $u(a)$ denote the level of electricity generated by the NPP corresponding to availability decision a. Let p_t denote the price of electricity at time t. Then $\mu(a_t, x_t, \phi)$ is given by

$$\mu(a_t, x_t, \phi) = \begin{cases} -\phi_c & \text{if } a_t = 1 \\ -c_r(x_t, \phi_r) & \text{if } a_t = 2 \\ p_t u(a_t) - c_o(x_t, a_t, \phi_o) & \text{if } a_t > 2, \end{cases} \quad (2.8)$$

where $c_r(x, \phi_r)$ is the expected cost of refueling in state $x, c_o(x, a, \phi_o)$ is the expected cost of operating a plant in state x at level a, and ϕ_c is the present value of costs associated with closing and decommissioning the NPP. Thus, the vector of unknown profit-function parameters is given by $\phi = (\phi_c, \phi_r, \phi_o)$.

Note that it is impossible to identify the location and scale of the utility's profit function using only data on operating histories. Therefore, we must impose an arbitrary normalization of location and scale. The location normalization can be imposed by assuming that $\mu(a, x, \phi) = 0$ for a prespecified state and decision pair (a, x). In our case, we follow Rust and Rothwell (1995) and normalize the present value of decommissioning costs to be 0; that is, $\phi_c = 0$. To simplify our model, we assume that the price of electricity varies only over the season but not across years, a reasonable assumption given the relative constancy of the price of electricity over the past decade and the slow 1–2% projected growth rate in demand to 2010 (Office of Technology Assessment 1993, p. 76). Under this assumption it is convenient to normalize the profit function by dividing π by the product of the plant's size and the electricity price p. The scale normalization is completed by assuming that the normalized error term ε_t has a standard Type I extreme-value distribution. By normalizing this way, we avoid the need to carry the electricity price and the plant's size as additional state variables in the DP model. Although this normalization reduces the computational burden of solving the DP model, it entails the implicit assumption that the optimal strategy for operating a plant is independent of its size. We plan to relax and test this assumption in future work.

We also follow Rust and Rothwell (1995) in making the standard simplifying assumption that the transition density λ can be factored as

$$\lambda(x_{t+1}, \varepsilon_{t+1} | x_t, \varepsilon_t, a_t) = p(x_{t+1} | x_t, a_t, \psi) q(\varepsilon_{t+1}), \quad (2.9)$$

where ψ is a vector of unknown parameters characterizing the transition density for the observable part of the state and control variables. Equation (2.9) is known as a "conditional independence assumption" because it implies that ε_{t+1} is independent of ε_t conditional on (x_t, a_t). Under the additional assumption that the marginal distribution of ε_t is Type I extreme value, Rust (1988, 1995) showed that the conditional choice probabilities, $P_t(a|x)$, are given by the classical multinomial logit formula

$$\begin{aligned} P_t(a|x) &= \int I\{a = \alpha_t(x, \varepsilon)\} q(d\varepsilon) \\ &= \frac{\exp\{v_t(x, a)\}}{\sum_{a' \in A_t(x)} \exp\{v_t(x, a')\}}, \end{aligned} \quad (2.10)$$

where the v_t are *expected value functions* given by the recursion formula

$$v_t(x, a) = \mu(x, a, \phi) + \beta \int \log \left[\sum_{a' \in A_t(x')} \exp\{v_{t+1}(x', a')\} \right] p(dx' | x, a, \psi). \quad (2.11)$$

The v_t functions given in (2.11) are related to the value function $V_t(x, \varepsilon)$ by the identity

$$V_t(x, \varepsilon) = \max_{a \in A_t(x)} [v_t(x, a) + \varepsilon(a)], \quad (2.12)$$

where we recall that the set $A_t(x)$ represents the set of feasible actions available to the operator in state x at time t.

We compute the solution to the DP model by backward induction using the recursion equation (2.11). The implied stochastic process for the observed state and control variables $\{x_t, a_t\}$ constitutes the DP model's prediction of a plant's optimal operating strategy. These predictions, however, depend on a vector of unknown parameters, $\theta = (\beta, \phi, \psi)$, specifying the discount factor, the unknown parameters of the profit function, and the law of motion for the state variables. We can estimate θ by maximum likelihood as follows. The Graybook data provide observations on the realization of the observed state and control variables for the sample of U.S. NPP's that were operational at the start of our sample in January 1989. Denote this data by $\{x_t^i, a_t^i\}, t = \underline{t}_i, \ldots, \bar{t}_i, i = 1, \ldots, N$. Given the conditional choice probability $P_t(a|x)$ in Equation (2.10) and the decomposition of the transition density $\lambda(x', \varepsilon' | x, \varepsilon, a)$ in Equation (2.9), it is straightforward to estimate the unknown parameter vector $\theta = (\beta, \phi, \psi)$ by maximum likelihood using the (full) likelihood function

$$L_f(\theta) = \sum_{i=1}^{N} \sum_{t=\underline{t}_i+1}^{\bar{t}_i} \log[P(a_t^i | x_t^i, \theta) p(x_t^i | x_{t-1}^i, a_{t-1}^i, \psi)]. \quad (2.13)$$

In practice, we estimate θ in a two-stage process: ψ is estimated from the partial likelihood function $L_1(\psi)$ given by

$$L_1(\psi) = \sum_{i=1}^{N} \sum_{t=\underline{t}_i+1}^{\bar{t}_i} \log[p(x_t^i | x_{t-1}^i, a_{t-1}^i, \psi)], \quad (2.14)$$

and the remaining parameters are estimated from the partial likelihood function $L_2(\beta, \phi | \hat{\psi})$ given by

$$L_2(\beta, \phi | \hat{\psi}) = \sum_{i=1}^{N} \sum_{t=\underline{t}_i+1}^{\bar{t}_i} \log[P(a_t^i | x_t^i, \beta, \phi, \hat{\psi})]. \quad (2.15)$$

Rust (1988) established the consistency and asymptotic normality of the two-stage and full maximum likelihood estimators. The covariance matrix for the parameters (β, ϕ) will not be consistently estimated from the second-stage partial

Table 2. Definitions of ϕ Parameters

Parameter	Description
β	Monthly discount factor (fixed at β = .999)
$\phi_{a=1}$	Expected present discounted value of costs of decommissioning NPP (normalized to 0)
$\phi_{a=2}$	Expected cost of refueling the NPP
$\phi_{a=2,f=2}$	Extra refueling cost following a forced outage
$\phi_{a=3,f=3}$	Per period cost of NPP shutdown during a major problem spell
$\phi_{a=3,f<3}$	Cost of NPP shutdown during an operating spell
$\phi_{a=3,f=2}$	Extra cost of shutdown due to forced outage
$\phi_{d,u>0}$	Effect of operating-cycle duration on expected profits (given positive availability)
$\phi_{u\in(u(a_t),u(a_{t+1})]}$	Expected profit of operating the NPP at an availability factor in the interval $(u(a_t),u(a_{t+1})]$
$\phi_{u=1}$	Expected profit of operating the NPP at 100% availability
$\phi_{u=1,f=2}$	Reduction in profits from operating NPP at 100% availability when there is a forced outage signal
$\phi_{\text{dec}},\ldots$	Adjustment to profit for an outage in December, January, etc.

likelihood from estimation noise in the first-stage parameters $\hat{\psi}$. We used the two-stage estimates of θ as a starting point for maximization of the full likelihood function $L_f(\theta)$, yielding consistent estimates of the covariance matrix and fully efficient estimates of θ.

3. ESTIMATION OF THE DP MODEL

This section presents structural estimation results for the parameters of the profit function and the parameters of the law of motion for the observed state variables of the DP model introduced in Section 2. We have estimated an unrestricted version of the profit function $\mu(a,r,d,f,\phi)$ defined by a 23×1 vector of coefficients defined in Table 2. We refer to the profit function as unrestricted because we are estimating coefficients representing normalized profits for various combinations of (x,a) rather than specifying the cost-function parameters and writing profits as electricity revenues minus costs as in Formula (2.8). Furthermore, the profit-function estimates are unrestricted in the sense that we do not impose any monotonicity conditions on the cost function—for example, that costs of running at higher availability are greater than the costs of running at lower availability and the cost of a refueling shutdown is greater than the cost of a shutdown for ordinary mid-cycle maintenance, and so forth.

As we noted earlier, we cannot identify the level of profits from the NRC data on NPP operations alone, so we made identifying normalizations of the location and scale of profits as described at the end of Section 2. Note that the unrestricted specification also includes monthly dummies to reflect seasonal variations in the opportunity cost of an outage. Because the 12 monthly dummies add up to a constant term, we also imposed an additional identifying normalization that the dummy coefficient for November is 0. Table 3 presents the definitions of the individual components of the ψ coefficients—the unknown coefficients characterizing the law of motion $p(x_{t+1}|x_t,a_t,\psi)$ described in Section 2.

Tables 4 and 5 present the maximum likelihood estimates of the ϕ and ψ parameters of the DP model over the 1989–1994 sample period using the full likelihood function $L_f(\theta) = L_f(\beta,\phi,\psi)$ given in Equation (2.13). The estimates are conditioned on a fixed value of the discount factor, $\beta = .999$, corresponding to a real annual interest rate of 1.2%. As in previous work, it is difficult to identify the exact value of β precisely, although a likelihood ratio test decisively rejects the hypothesis that operators use discount rates higher than 10% per year. The results tell us that operators seem to be extremely farsighted in their decision making, and the fact that the likelihood function is essentially flat for discount factors greater than $\beta = .999$ may indicate that the operator's objective is simply to maximize total (undiscounted) expected profits over the remaining lifetime of the NPP.

The parameter estimates in Table 4 are similar to those of Rust and Rothwell (1995), so we refer the reader to that article for a detailed discussion of their interpretation so that we can concentrate on describing the new findings we have obtained for our new sample and our new three-spell version of the DP model. Recall that the magnitudes of the coefficient estimates are not meaningful because of our normalizations of the location and scale of the profit function. The relative values of the various coefficient estimates, however, are meaningful in this model. Starting with the coefficient of the profit function corresponding to 100% availability under "normal" conditions, $\hat{\phi}_{u=1} = 1.548$, we

Table 3. Definitions of ψ Parameters

Parameter	Description
$\psi_{rf}(1)$	Constant term in probability of forced outage following refueling spell
$\psi_{rf}(t)$	Age term in probability of forced outage following refueling spell
$\psi_{ro}(d_t = i)$	Dummy on the ith month in refueling-spell-duration model, i = 1,2,3,4
$\psi_{ro}(d_t \geq 5)$	Dummy on fifth and later months in refueling-spell-duration model
$\psi_{ro}((d_t - 4)(d_t \geq 5))$	Slope term for fifth and later months in refueling-spell-duration model
$\psi_{of}(1)$	Constant term for probability of forced outage during an operating spell
$\psi_{of}(t)$	Age term for probability of forced outage during an operating spell
$\psi_{of}(d_t)$	Duration term for probability of forced outage during an operating spell
$\psi_{of}(d_t^2)$	Squared duration term for probability of forced outage during an operating spell
$\psi_{of}(f_t = 1)$	Dummy for forced outage in month t on probability of forced outage in month $t + 1$
$\psi_{om}(1)$	Constant term for probability of major problem during an operating spell
$\psi_{om}(t)$	Age term for probability of major problem during an operating spell
$\psi_{om}(d_t)$	Duration term for probability of major problem during an operating spell
$\psi_{om}(f_t = 1)$	Dummy for forced outage in month t on probability of major problem in month $t + 1$
$\psi_{mo}(1)$	Constant term in probability of coming out of a major problem spell
$\psi_{mo}(d_t)$	Duration term in probability of coming out of a major problem spell

Table 4. Full Information Maximum Likelihood Estimates of ϕ Parameters

Parameter	Estimate	Standard error	t statistic
$\phi_{a=2}$	−2.636	.341	−7.72
$\phi_{a=2,f=2}$	−3.704	.295	−12.57
$\phi_{a=3,f=3}$	−1.768	.667	−2.65
$\phi_{a=3,f<3}$	−2.220	.194	−11.46
$\phi_{a=3,f=2}$	−4.408	.331	−13.31
$\phi_{d,u>0}$	$-.703 \times 10^{-1}$	$.145 \times 10^{-1}$	−4.84
$\phi_{u\in(0,.25]}$	−3.527	.199	−17.74
$\phi_{u\in(.25,.5]}$	−3.104	.188	−16.51
$\phi_{u\in(.5,.75]}$	−2.194	.177	−12.40
$\phi_{u\in(.75,1)}$	−1.047	.171	−6.11
$\phi_{u=1}$	1.548	.168	9.24
$\phi_{u=1,f=2}$	−5.934	.195	−30.41
ϕ_{dec}	−.692	.291	−2.38
ϕ_{jan}	−.636	.260	−2.45
ϕ_{feb}	−.743	.231	−3.21
ϕ_{mar}	.353	.223	1.58
ϕ_{apr}	.336	.237	1.42
ϕ_{may}	.143	.239	.60
ϕ_{june}	−.521	.254	−2.05
ϕ_{july}	−.694	.255	−2.72
ϕ_{aug}	−1.551	.242	−6.41
ϕ_{sep}	.210	.243	.87
ϕ_{oct}	.241	.288	.84

NOTE: $\log(L_f(\hat{\phi}, \hat{\psi})) = -9{,}011.66$; $N = 7{,}526$.

see that profits corresponding to successively lower levels of availability, $\hat{\phi}_{u\in(.75,1)}, \ldots, \hat{\phi}_{u\in(0,.25]}$, decrease monotonically, as expected. The only anomaly is that the profit corresponding to a complete shutdown of the NPP for the month, $\hat{\phi}_{a=3,f<3} = -2.22$, is lower than the profit corresponding to running the NPP at availability factors between 1% and 75%. This is similar to the finding of Rust and Rothwell (1995), and we believe that the explanation for the anomaly is the same: Because our specification of the DP model does not fully capture the regenerative, investment value of a mid-cycle preventive maintenance outage, the coefficient estimate of $\hat{\phi}_{a=3,f<3} = -2.22$ might be capturing both the current costs of such an outage and the present value of reduced future costs. This would lead the coefficient estimate to have an upward bias and could explain the counterintuitive implication that a complete shutdown of the NPP is estimated to be more profitable per period than running the NPP at partial availability. Similar comments apply to the estimated value of the monthly cost of refueling the NPP, $\hat{\phi}_{a=2}$. The only way our DP model is able to reflect the regeneration that occurs from preventive maintenance during a refueling outage is to reset the spell duration counter d_t back to 1 [see Eq. (2.4)]. We acknowledge that our estimate of $\hat{\phi}_{a=2}$ might be biased upward if this simplified way of modeling the regenerative effect of refueling outages does not fully capture the true regenerative effect of the refueling. The bias equals the expected present discounted value of the reduction in operating costs in future operating spells—that is, the difference between realized O&M costs and the higher levels that would have been incurred had the preventive maintenance activities not been undertaken during the current refueling spell.

Most of the other parameter estimates appear economically plausible. One of the most interesting findings is that NPP operators are highly averse to "imprudent" operation of their NPP's in the sense that they are very reluctant to run the NPP at 100% availability after receiving a forced outage signal. The expected loss of such imprudent behavior is $-4.4 = \hat{\phi}_{u=1} + \hat{\phi}_{u=1,f=2}$. This is significantly greater than the loss that would be incurred by running the NPP at lower availability levels. Once again, this coefficient estimate could be biased because it reflects the expected present value of costs of repairing damage to the reactor that might have otherwise been incurred had the operator ignored the forced outage signal and insisted on running the NPP at 100% availability. These costs also could include "goodwill costs," such as fines by the NRC.

The expected loss incurred in each month of a major problem spell, $\hat{\phi}_{a=3,f=3} = -1.77$, is estimated to be significantly lower than the loss incurred during a refueling outage or during a mid-cycle preventive-maintenance outage. This also could be an anomalous finding because expensive capital upgrades and retrofits typically occur during major problem spells (e.g., replacement of steam generators). For reasons already discussed, we expect that this coefficient could be biased upward by the difference in the present value of O&M costs following the major problem spell and the levels of O&M costs that would have been incurred if the capital upgrades and other repairs had not been undertaken.

In view of these problems, one cannot interpret the individual profit-function estimates in Table 4 too literally. Few of them can be viewed as representing expected profits in the current month: Instead, most of them reflect a combination of current profits and expected present values of future profits. This is not problematic for our analysis because our interest focuses on recovering the value function V rather than the current profit function μ. We argue that our unrestricted specification of μ provides the flexibility necessary to approximate V even though we have not fully specified

Table 5. Full Information Maximum Likelihood Estimates of ψ Parameters

Parameter	Estimate	Standard error	t statistic
$\psi_{rf}(1)$	−.847	.223	−3.80
$\psi_{rf}(t)$	$.107 \times 10^{-2}$	$.129 \times 10^{-2}$	.83
$\psi_{ro}(d_t = 1)$	−5.928	1.004	−5.91
$\psi_{ro}(d_t = 2)$	−1.750	.145	−12.10
$\psi_{ro}(d_t = 3)$	$.132 \times 10^{-1}$	.112	.12
$\psi_{ro}(d_t = 4)$	.474	.162	2.93
$\psi_{ro}(d_t \geq 5)$	−.137	.408	−.33
$\psi_{ro}((d_t - 4)(d_t \geq 5))$	.267	.226	1.18
$\psi_{of}(1)$	−1.444	.106	−13.58
$\psi_{of}(t)$	$.441 \times 10^{-3}$	$.352 \times 10^{-3}$	−1.25
$\psi_{of}(d_t)$	$.320 \times 10^{-1}$	$.222 \times 10^{-1}$	1.44
$\psi_{of}(d_t^2)$	$-.210 \times 10^{-2}$	$.124 \times 10^{-2}$	−1.70
$\psi_{of}(f_t = 1)$	.340	$.755 \times 10^{-1}$	4.50
$\psi_{om}(1)$	−6.905	.762	−9.06
$\psi_{om}(t)$	$-.548 \times 10^{-2}$	$.176 \times 10^{-2}$	−3.12
$\psi_{om}(d_t)$	.166	$.392 \times 10^{-1}$	4.24
$\psi_{om}(f_t = 1)$	1.259	.479	2.63
$\psi_{mo}(1)$	−3.350	.556	−6.02
$\psi_{mo}(d_t)$	$.499 \times 10^{-1}$	$.307 \times 10^{-1}$	1.62

NOTE: $\log(L_f(\hat{\phi}, \hat{\phi})) = -9{,}011.66$; $N = 7{,}526$.

Table 6. Predicted Versus Actual Choice Probabilities: Full Sample

Action	NP	DP
a = 1 (close)	5.3×10^{-4}	5.2×10^{-4}
a = 2 (refuel)	.17154	.16539
a = 3 (u = 0%)	.07667	.07639
a = 4 (u = 13%)	.01236	.01242
a = 5 (u = 38%)	.01887	.01896
a = 6 (u = 63%)	.04690	.04714
a = 7 (u = 88%)	.14762	.14836
a = 8 (u = 100%)	.52551	.53083

NOTE: $\chi^2(7)$ = 53.19.

the dynamic structure of how current maintenance and capital upgrade investments influence future O&M costs.

The remaining coefficients in Table 4 include the effect of gradual deterioration within each operating cycle, captured by the negative coefficient estimate on the duration term $\hat{\phi}_{d,u>0}$ and seasonal variations in the demand for power (seasonal variations in the price of electricity) as reflected in the estimated monthly dummy variables. As expected, the estimation results reveal that the most costly time to shut down a plant is in the winter or summer when power demand is at its peak, and the best times to shut down the NPP for maintenance or refueling is during the early spring (March or April) or early fall (September or October).

The estimation results in Table 4 are based on a horizon of $T = 480$ months. In effect, we are assuming that NPP operators were convinced that the maximum license span would always be 40 years and that they would never have an opportunity to apply for a license extension. We tested this expectational hypothesis by reestimating the DP model using alternative horizons ranging from 40 to 60 years. We found that the estimated log-likelihood function declined monotonically as the assumed length of the license span increased from 40 years. In this sense, our empirical results seem to confirm the view that, even though the NRC introduced a procedure for applying for license extensions in 1991, most NPP operators were extremely pessimistic about the chances that their applications would actually be granted.

Table 5 presents the estimates of the ψ parameters characterizing the stochastic law of motion $p(x_{t+1}|x_t, a_t, \psi)$ for the observed state variables x_t. The estimates are consistent with the findings obtained by Rust and Rothwell (1995) for the sample period 1984 to 1993, except (a) the estimates imply that the refueling durations were significantly shorter on average after 1989 than they were from 1984 to 1988, (b) the level of serial correlation in forced outages has decreased, and (c) the risk of entering a major problem spell *decreases* with the age of the NPP, increases with duration of the operating spell, and is significantly higher following a forced outage. As discussed in Section 1, it is difficult to detect aging effects in our sample of NPP's. Older NPP's appear to have a systematically lower risk of forced outages and major problems than younger NPP's. Apparently the combined effects of technological progress and learning-by-doing outweigh the effects of age-related deterioration.

Table 6 presents the results of a chi-squared goodness-of-fit test, using the version of the statistic developed by Andrews (1988) that accounts for covariates appearing in the conditional choice probability $P(a|x,\hat{\theta})$. In general, there is a problem in comparing predicted and actual choices because there are 108,000 possible (x_t,t) cells, but in our dataset there are observations in only 6,131 distinct (x_t,t) cells with an average of only 1.22 observations per cell. To compare the predictions of the DP model to the data, we need to aggregate over cells. If X denotes a collection of (x_t,t) cells, we define the nonparametric (NP) and parametric (DP) estimates of the conditional probability $P(a|X) \equiv P(a|x \in X)$ as follows:

$$\begin{aligned}\hat{P}(a|X) &= \int_{x\in X} \hat{P}(a|x)\hat{F}(dx|X) \\ &\equiv \frac{1}{N}\sum_{i=1}^{N} I\{a_i = a, x_i \in X\} \quad \text{(NP)} \\ P(a|X,\hat{\theta}) &= \int_{x\in X} P(a|x,\hat{\theta})\hat{F}(dx|X) \\ &\equiv \frac{1}{N}\sum_{i=1}^{N} P(a|x_i,\hat{\theta})I\{x_i \in X\} \quad \text{(DP)}. \quad (3.1)\end{aligned}$$

Andrews's (1988) version of the chi-squared goodness-of-fit statistic allows one to choose various partitions $(X_1,\ldots,X_J)$ of the set of all possible (x_t,t) cells in the calculation of the statistic, which is simply a quadratic form measuring how close the $8\times J$ vectorized collection of residuals $[P(a_i|X_j,\hat{\theta}) - \hat{P}(a_i|X_j)]$ is to 0. The results in Table 6 are for the special case in which $J = 1$, so X_1 is a collection of all possible (x_t,t) cells, and the chi-squared statistic has an asymptotic chi-squared distribution with 7 df.

Although the predicted and actual choice probabilities seem quite close to each other, the magnitude of the chi-squared statistic indicates a decisive rejection of the DP model. Reasons the DP model is rejected include (a) NPP operators might not be discounted profit maximizers, (b) there might be plant-level heterogeneity that is not accounted for by the DP model, and (c) the conditional independence assumption [see Eq. (2.9)] might not be valid. Unfortunately the chi-squared statistic by itself is not informative as to which of the assumptions underlying the DP model are most likely to be violated.

Although rejected by the chi-squared test, we believe that the estimated DP model is able to do a reasonable job of predicting the aggregate behavior of the nuclear industry. Figure 1 plots the predicted versus the actual distributions of the durations of operating spells. The DP model correctly predicts that the *planned* duration of operating spells is 18 months. Due to uncertain events, however, the *realized* durations of actual operating spells tend to be shorter, with an average duration of 14.61 months. We can see from Figure 1 that the DP model predicts that realized operating-spell durations have a mean of 14.58 months, which is very close to what we observe. The fact that the DP model is able to closely mimic the shift in both the planned and realized durations of operating spells following the TMI accident (see Rust and Rothwell 1995 for details) provides indirect

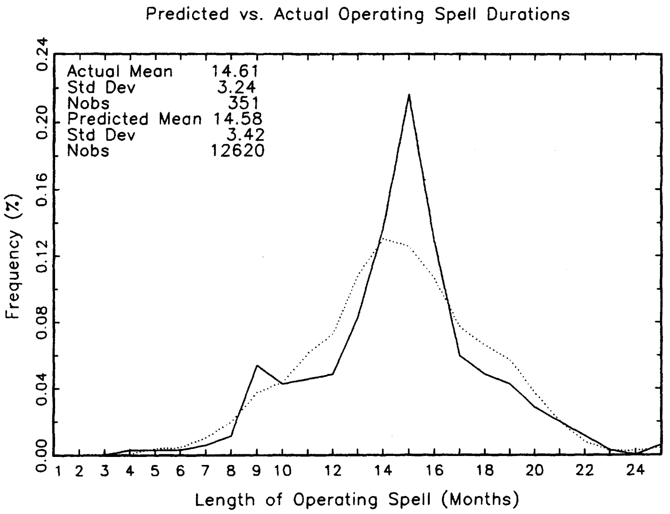

Figure 1. Predicted Versus Actual Distribution of Operating Spell Durations: ———, Actual, · · · ·, Simulated.

evidence of its usefulness for forecasting and policy simulation.

We conclude this section with Figure 2, which provides evidence of the DP model's ability to make accurate out-of-sample forecasts of the nuclear industry as a whole. We ran 20 stochastic simulations of the estimated DP model using the state of the nuclear power industry in January 1984 as initial condition. (If a plant came online after January 1984 we began its simulation from the first month that it became operational.) We refer to these as "out-of-sample" simulations because the DP model was estimated on data from January 1989 to December 1994 and the simulations are run over the entire 11-year span from January 1984 to December 1994. We feel that this is a good test of the ability of the DP model to make long-run forecasts because the divergence between predicted and actual trajectories usually increases with time in complicated nonlinear stochastic models.

The top panel of Figure 2 plots the simulated versus the actual number of operating NPP's. The mean number of plants in the 20 industry simulations tracks the actual number of plants remarkably closely. To get a sense of the potential variance in the simulations, we plotted 90% confidence bands (computed pointwise as the highest and lowest number of firms observed in the 20 simulations at each point in time). The fact that these bands are relatively tight suggests that published predictions of a rash of "early retirements" are way off target. The bottom panel of Figure 2 plots simulated versus actual industry generating capacity. In this case the mean capacity in the 20 simulations slightly underpredicts actual capacity levels and the actual trajectory for industry capacity virtually coincides with the upper confidence band. The explanation for the apparent discrepancy between the top and bottom panels of Figure 2 is that the six plants that were closed since 1986 happened to be much smaller than the average size NPP.

4. PREDICTING THE LIFETIME OF THE NUCLEAR INDUSTRY UNDER ALTERNATIVE LICENSING RULES

In this section we use the estimated DP model to predict the evolution of the nuclear power industry under two policies regarding operating license extensions, (1) *no license extensions,* where each NPP has a maximum license term of 40 years, and (2) *costless 20-year license extensions,* where each NPP has a maximum license term of 60 years. The policy that will finally emerge from the NRC's deliberations will likely end up somewhere between these two extremes. In particular, if the policy imposed large costs and uncertainties on the applicant, as discussed in Section 1, it would be necessary to explicitly model the decision of whether to apply for a license extension as part of the overall DP problem. The solution to this problem will depend on the operator's perceptions of the expected costs of submitting an application and the likelihood that the application will be successful. Thus, the two policy alternatives we consider correspond to cases of infinite and zero applications costs, respectively, and subjective probabilities of 0 and 1 that the application will be accepted.

We begin by forecasting the evolution of the nuclear power industry under the assumption of 40-year operating licenses. Our projections depend on many assumptions

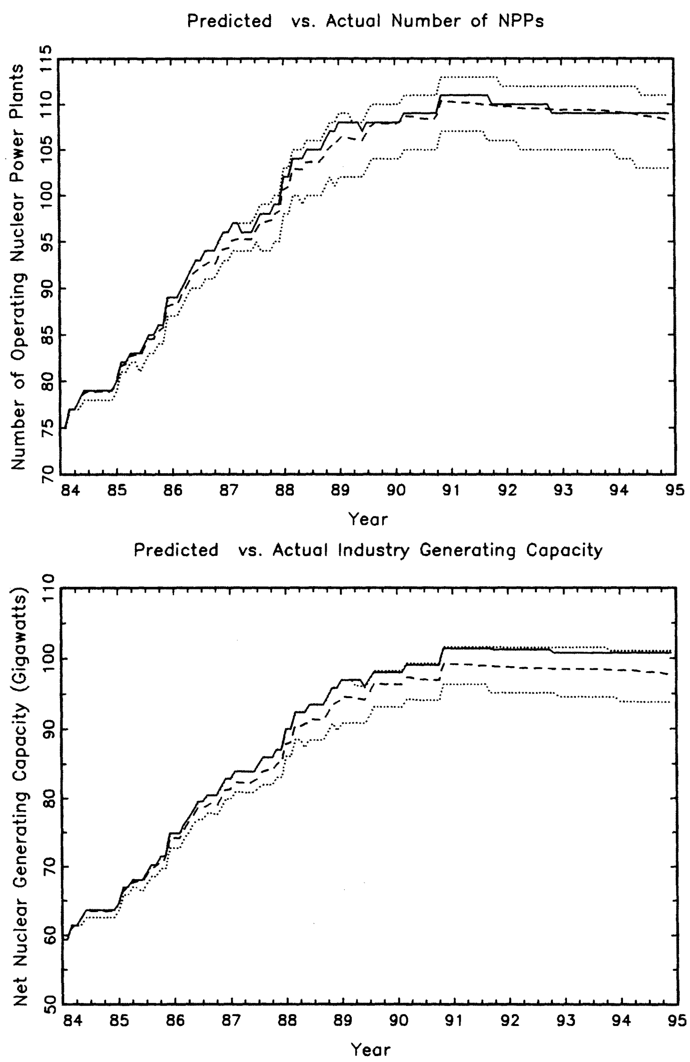

Figure 2. Comparison of Actual Versus Simulated Number of Plants and Total Generating Capacity, 1984–1995: ———, Actual; - - -, Simulation Mean; · · · · (upper), 5th Percentile; · · · · (lower), 95th Percentile.

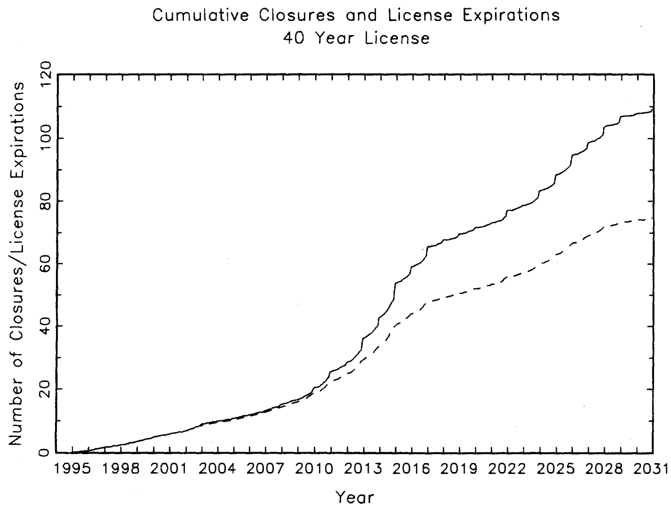

Figure 3. Simulated Closures Under 40-Year Operating Licenses: ———, Exits, - - -, Closures Prior to License Expiration.

about the regulatory and economic environment, including (a) the NRC and PUC regulatory regimes will not change, (b) all utilities sell electricity at a common price over a competitive power grid, (c) the real price of electricity will remain at current levels, (d) the government will not impose carbon taxes or costly environmental regulations on fossil fuels, (e) real NPP O&M costs will remain at current levels, (f) the Department of Energy will be successful in developing a long-term waste repository for spent nuclear fuel, (g) there will be no other major changes in the expected costs of NPP decommissioning, and (h) there is no sudden increase in the rates of deterioration in major plant and reactor components for the first 60 years of the NPP's life.

Figure 3 simulates the evolution of the nuclear power industry under these assumptions and the assumption of 40-year operating licenses. We forecast that the last NPP will close in 2031 when the youngest currently operating NPP (which came online in 1991) reaches the end of its 40-year license. Figure 3 plots the expected cumulative number of exits and NPP closures over time. The two curves coincide up to the year 2010 because all exits over this period are from "early retirements." We predict that a total of 20 NPP closures will occur between now and 2010. Note, however, that we predict that only three NPP's will be closed between 1995 and 2000. This result is in marked contrast to other forecasts, such as the prediction quoted in the introduction, that as many as 25 NPP's will be closed in the next few years. By the time the last NPP closes in 2031, we predict that a total of 75 NPP's will have closed before their license expiration dates, leaving only 32 NPP's to be operated for the full duration of their operating licenses.

Figure 4 presents forecasts of the evolution of the nuclear power industry under a 60-year license span and compares the DP model's prediction of the number of firms and industry generating capacity under 40- and 60-year licenses. Not surprisingly, the 20-year license extension also extends the life of the nuclear industry by 20 years, from 2031 to 2051. It is more surprising to note, however, that under 60-year licenses there is virtually no drop-off in industry generating capacity until after 2031. Although closures peak in the year before the license expiration date, fully 22% of all NPP's are closed more than 10 years "prematurely" under the regime of 40-year operating licenses. Under 60-year operating licenses, however, only 3% of the NPP's close early. The reason for the significantly lower rate of early retirements under a 60-year license is a pure horizon effect: A longer license span gives the owner a longer horizon over which to recoup losses made during major problem spells or to pay back the capital investments made to extend plant longevity.

This "horizon effect" is verified in Figure 5, which plots the value functions under 40- and 60-year operating licenses. The highest curve in each diagram is the value corresponding to running the NPP at 100% availability when there are no forced outage signals. The quadratic shape of these curves is from the economics of the problem rather than an a priori functional-form specification because there are no age terms entering the per-period profit function (see the definition of the ϕ coefficients in Table 2). Because the probability of major problems and forced outages decreases with the plant age, NPP's become steadily more profitable per unit of time as they age. This fact is responsible for the initial upward slope in the value functions. There is also a "horizon effect" that forces value functions to decrease after some point, however: As the NPP approaches its license-expiration date, expected discounted profits fall. Thus, the quadratic shape of the value function is the result of the trade-off between the increase in NPP profitability with age and the horizon effect.

Figure 5 shows that, under "normal" circumstances (i.e., when running at 100% with no forced-outage signal, or even when the NPP is shut down for repair following a forced-outage signal), there is virtually no chance that the NPP will be closed. Both panels show, however, that once the NPP enters a major problem spell the level of expected discounted profits falls dramatically and the chances that the NPP will be closed increase substantially. Under a 40-year license the optimal policy is to close the plant if a major

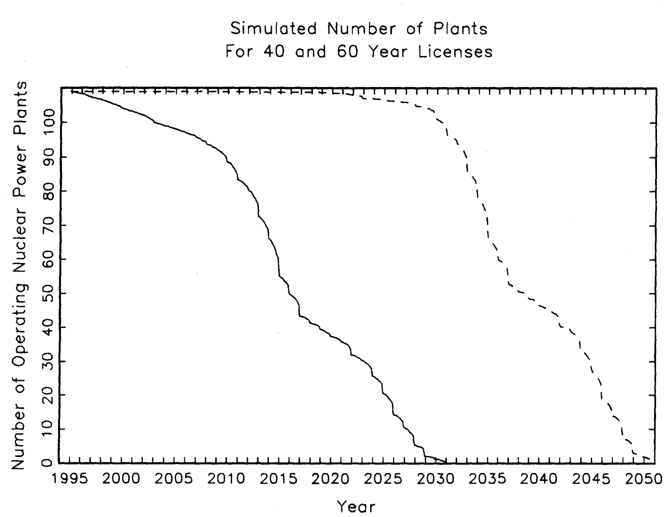

Figure 4. Predicted Number of NPP's Under 40- and 60-Year Operating Licenses: ———, 40-Year License; - - -, 60-Year License.

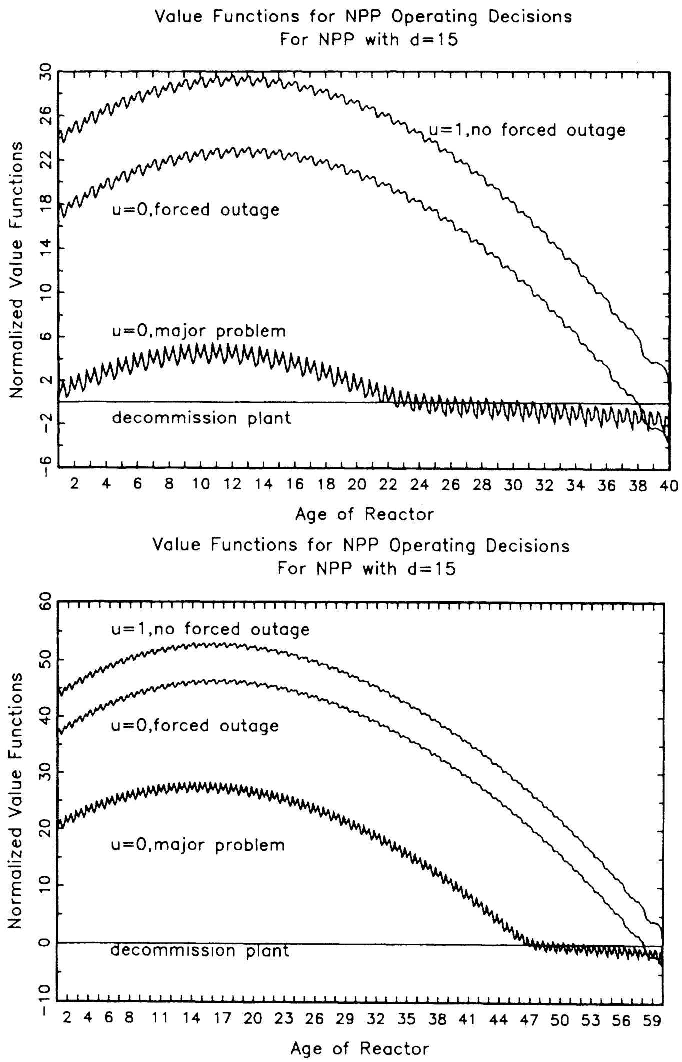

Figure 5. Value Functions Under 40- and 60-Year Operating Licenses.

problem signal is received in the first year of its operating license or if it is received by a plant that has fewer than 16 years remaining in its license span. Under a 60-year license, however, it will be optimal to incur the costs of a major problem spell in all but the final 13 years of the operating license. In both cases it is clear that NPP operators are making optimal decisions about plant life with the remaining horizon in mind: The longer the remaining horizon, the greater the chances that major investments incurred during a major problem spell can be recouped. This result suggests that, had a policy of 60-year operating licenses been in effect at the time of the TMI accident in 1979, far fewer new NPP projects would have been cancelled. The reason is that industrywide NRC-mandated administrative shutdowns in the wake of the TMI accident can be regarded as major problem spells. Under 40-year operating licenses, Figure 5 shows that, if a major problem occurs in the first year of operation of a new NPP, the operator would be nearly indifferent between continuing to operate the plant and closing it. Under 60-year operating licenses, however, Figure 5 shows that operators would unambiguously prefer the option of keeping the plant open rather than closing it down.

Although the levels of value functions are not meaningful because of our arbitrary identifying scale and location normalizations, the ratio of the value function for a 60-year license span to the value function under a 40-year license span is a meaningful quantity. One can see from Figure 5 that moving to a 60-year license doubles the value function. Thus, we conclude that a 60-year license span would double the (discounted) profits of the nuclear power industry. Our simulations of net nuclear electricity generation (not shown due to limited space) show that a costless 20-year license extension would also double the total (undiscounted) electricity production over the remaining life of the industry.

5. CONCLUDING REMARKS

The policy simulations conducted in Section 4 are subject to three important qualifications. First, the accuracy of the predictions depends on the assumption that there will be no major changes in the current economic environment or regulatory regime. Deregulation of electricity generation could lead to significant changes in the price of electricity and dramatically affect the relative profitability of nuclear power plants. Second, the predictions were based on the assumption that there is no critical aging threshold—that is, an age beyond which there is a sudden acceleration in the rate of deterioration in major reactor components. Our estimates assume that regular maintenance and capital upgrades are sufficient to keep the reactor running safely. If this is not the case, the gains to adopting a 60-year license span might be much less than we have predicted. Third, the accuracy of our predictions depends on the accuracy of the DP model as a positive model of NPP operations. The chi-squared goodness-of-fit test statistic in Section 3 indicates that our model is misspecified. One source of specification error is the homogeneity assumption—that is, that plants differ only by observable state variables. Differences in economic performance could result in more plant closures than predicted by our simple homogeneous specification of the DP model. It is well known that there are wide differences among plants in their reliability and levels of O&M costs. If these differences are because certain utilities operate "lemon" reactors, we may be seeing more plant closures than predicted by our simple homogeneous specification of the DP model. If the differences in performance are a result of the fact that some management teams are simply not behaving optimally, however, then a fundamental maintained assumption underlying the DP model is called into question. Our forecasts may be too pessimistic if deregulation leads "good managers" to buy poorly performing NPP's owned by "bad managers" and dramatically improve their reliability and cost-effectiveness.

We conclude by reemphasizing that this article has focused on the question of the optimal lifetime of an NPP from the private perspective of the plant owner rather than the social perspective of the regulator. For reasons already noted in Section 1, we do not take a position on whether it is socially optimal to switch to a regime that permits costless 20-year extensions of NPP operating licenses. We believe, however, that our analysis provides a starting point for addressing this question.

ACKNOWLEDGMENTS

We are grateful for research support from the National Science Foundation, which provided central processing unit time on the Pittsburgh Supercomputer Center's C-90 supercomputer, and from IBM. We are also grateful to the Computer Capabilities for Regulatory Research project at the Center for Economic Policy Research, Stanford University. John Rust thanks the Bradley Foundation, whose generous grant for a project on Social Security also provided indirect support for this project. We would like to extend special thanks to helpful comments from workshop participants at New York University and Northwestern University and from J. Altonji, W. Brainard, M. Buchinsky, J. Cummins, I. Domowitz, C. Flinn, J. Hewlett, B. Meyer, W. Nordhaus, M. Rothwell, C. Taber, and an anonymous referee.

[Received January 1995. Revised August 1996.]

REFERENCES

Andrews, D. W. K. (1988), "Chi-Square Diagnostic Tests for Econometric Models: Theory," *Econometrica*, 56, 1419–1454.

Che, Y. K., and Rothwell, G. S. (1995), "Performance-Based Pricing at U.S. Electric Utilities," *The Energy Journal*, 16, 57–77.

Energy Information Administration (1983), *Nuclear Plant Cancellations: Causes, Costs, and Consequences*, Washington, DC: U.S. Department of Energy (DOE/EIA-0392).

——— (1995), *An Analysis of Nuclear Plant Operating Costs: A 1995 Update*, Washington, DC: U.S. Department of Energy (SR/OIAF/95-01).

Forest, L., Jr., Deutsch, T. R., and Schenler, W. W. (1988), "Cost Savings From Extended Life Nuclear Plants," in *Nuclear Plant Life Extension, Proceedings of American Nuclear Society Meeting at Snowbird, Utah* (Vol. 1), pp. 225–231.

Hewlett, J. G. (1991), "A Cost/Benefit Perspective of Extended Unit Service as a Decommissioning Alternative," *The Energy Journal*, 12 (Special Issue on Nuclear Decommissioning Economics), 255–272.

Institute for Nuclear Power Operations (1993), *1992 Performance Indicators for the U.S. Nuclear Utility Industry*, Atlanta: Author.

Lester, R. K., and McCabe, M. J. (1993), "The Effect of Industrial Structure on Learning by Doing in Nuclear Power Plant Operation," *The RAND Journal of Economics*, 24, 418–438.

Nuclear Regulatory Commission (1989–1994), *Licensed Operating Reactors: Status Summary Report* Washington, DC: Author (NUREG-0020).

Office of Technology Assessment, U.S. Congress (1993), *Aging Nuclear Power Plants: Managing Plant Life and Decommissioning*, Washington, DC: U.S. Government Printing Office (OTA-E-575).

Pasqualetti, M. J., and Rothwell, G. S. (eds.) (1991), "Special Issue on Nuclear Decommissioning," *The Energy Journal*, 12.

Rothwell, G. S. (1990), "Utilization and Service: Decomposing Nuclear Reactor Capacity Factors," *Resources and Energy*, 12, 215–229.

——— (1996), "Organizational Structure and Expected Output for Nuclear Power Plants," *Review of Economics and Statistics*, 56, 482–488.

Rothwell, G. S., and Rust, J. (1995), "A Dynamic Programming Model of U.S. Nuclear Power Plant Operations" Discussion Paper 410, Stanford University, Center for Economic Policy Research.

Rust, J. (1987), "Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher," *Econometrica*, 55, 999–1033.

——— (1988), "Maximum Likelihood Estimation of Discrete Choice Processes," *SIAM Journal of Control and Optimization*, 26, 1006–1024.

——— (1995), "Structural Estimation of Markov Decision Processes," in *Handbook of Econometrics, Vol 4*, eds. R. Engle and D. McFadden, New York: North-Holland, pp. 3082–3143.

Rust, J., and Rothwell, G. S. (1995), "Optimal Response to Shift in Regulatory Regime: The Case of the U.S. Nuclear Power Industry," *Journal of Applied Econometrics*, 10, S75–S118.

Shah, V. N., and MacDonald, P. E. (eds.) (1993), *Aging and Life Extension of Major Light Water Reactor Components*, Amsterdam: Elsevier.