On the Optimal Lifetime of Nuclear Power Plants

Geoffrey Rothwell
Stanford University
rothwell@leland.stanford.edu

John Rust
University of Wisconsin
jrust@thor.econ.wisc.edu

December 1995

Abstract: This paper presents a dynamic programming (DP) model of operating a nuclear power plant (NPP). In each period the operator decides whether to run the reactor, to shut it down for preventive maintenance or refueling, or permanently close the plant for decommissioning. The maximum life of a NPP is determined by the length of the operating license issued by the Nuclear Regulatory Commission. The optimal lifetime of a NPP (from the private perspective of the plant owner, as opposed to the social perspective of the regulator) is the solution to a generalized optimal stopping problem: the operator closes a plant as soon as the expected discounted value of future operating profits (losses) falls below the costs of decommissioning. This model extends the DP model of NPP operations introduced in Rust and Rothwell (1996) by allowing for the occurrence of "major problem spells." The DP model predicts that under ordinary operating conditions it is unlikely that an NPP will be closed, but the probability of decommissioning increases substantially during a major problem spell. We compare the evolution of the nuclear power industry from 1984 to 1994 to stochastic simulations of our estimated DP model to show that our model provides accurate out-of-sample predictions of early retirements of NPPs. We use the estimated DP model to forecast nuclear power generation under two policy scenarios: (1) the current 40-year license span with no possibility of extension and (2) a "costless" extension in operating licenses to 60 years. Our simulations show that an immediate, costless extension in operating licenses to 60 years would double the expected present discounted value of profits to continued operation of US NPPs.[1]

JEL Classification: C41–Duration Analysis

Keywords: dynamic programming, electricity generation, nuclear power plants, reactor aging, life extension

---

[1] We are grateful for research support from the National Science Foundation, which provided cpu time on the Pittsburgh Supercomputer Center's C-90 supercomputer, and from IBM to the Computer Capabilities for Regulatory Research project at the Center for Economic Policy Research, Stanford University. John Rust is grateful to 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 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.

1. Introduction

Much attention has been paid to the decline of the nuclear power industry in the US following the Three Mile Island (TMI) accident in March 1979 (see, e.g., Campbell, 1988). 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 (NPPs) since 1978. Orders for more than 100 NPPs were cancelled, some involving plants that were nearly complete, resulting in losses of tens of billions of dollars in planning and construction costs (EIA, 1983). In a previous paper (Rust and Rothwell, 1996) we estimated that more than 90% of the discounted gains to continued operation of existing NPPs were eliminated in the stricter post-TMI regulatory regime.

Despite the reduction in profits, more than 100 NPPs continue to operate in the US, constituting a combined generating capacity of 100 gigawatts or about 15% of the country's total electrical capacity. In 1992 these NPPs produced 620 billion kilowatt-hours, or 22% of the nation's electricity supply (INPO, 1993). Although aggregate demand for electricity is projected to grow at moderate rates of approximately 1-2% over the remainder of the 1990's, there will be a significant loss of generating capacity over the next three decades from retirements of aging NPPs. Some studies (e.g., Forest, Deutsch and Schenler, 1988, and Makovich, Forest and Fletcher, 1988) have estimated that extending the life of existing NPPs by 20 years could result in large savings — about $450 billion in current dollars — reducing the electricity bills of the average family by more than $200 per year. In addition, increasing recognition of the adverse environmental impacts of fossil fuels, including acid rain and the "greenhouse effect," and national interest in having a secure and dependable long term source of energy has lead to a reconsideration of the "nuclear option" (see Cohen, 1990).

Under the Atomic Energy Act, the Nuclear Regulatory Commission (NRC) issues operating licenses for a maximum term of 40 years. Assuming there will be no new investment in NPPs, it is straightforward to predict that nuclear power is on a steady path to extinction in the US over the next 35 years as NPPs reach the end of their license lifetimes. In 1991 the NRC developed a preliminary set of procedures that would allow NPPs to apply for an extension of their operating licenses by an additional 20 years. Although the NRC is currently solidifying its procedures for extending operating licenses, the combination of increased operating costs and a variety of age-related problems (for example, radiation-induced embrittlement of NPP reactor vessels and corrosion-induced failures in steam generators) have lead to the presumption that license extension would not be economic. For example, more recent cost-benefit studies, such as Hewlett (1991a), concluded that unless there is a reduction in the high levels and growth rates in operating and maintenance costs experienced by NPPs 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). Thus utilities might not find it profitable to continue operating existing NPPs for the full 40-year duration of their initial license. Indeed, since 1988 six NPPs have been closed more than ten years before their license expiration dates: La

Crosse, Fort St. Vrain, Rancho Seco, Trojan, Yankee Rowe, and San Onofre 1. 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 paper presents an empirical model of an operator's decision whether to operate or close a plant. Our model extends a dynamic programming (DP) model of optimal NPP operation developed by Rust and Rothwell (1996) 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, or extended NRC-mandated forced 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 decommissioning increases substantially during a major problem spell. We compare the evolution of the nuclear power industry from 1984 to 1994 to stochastic simulations of our estimated DP model. We show that the model provides accurate out-of-sample predictions of early retirements of NPPs. Our simulations indicate that the prediction of 25 closures in the next few years is too pessimistic: we predict that only 3 additional NPPs will be closed between now and 2000, but 20 NPPs will be closed by 2010.

We use the estimated DP model to forecast nuclear power generation under two policy scenarios: (1) 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. Paradoxically, the DP model also predicts that it is optimal to close a plant that experiences a major problem in the first year of operation.[2] However if a major problem occurs when the NPP is between 2 and 24 years old the optimal strategy is to incur the large expenses to fix the problem and continue operating. Since most existing US NPPs are between 2 and 24 years old, this result could explain why more than 100 new NPP projects were cancelled in the wake of the TMI accident but only 5 older existing NPPs have been closed since 1988.

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. This suggests that if a 60-year license span had been in effect at the time of the TMI accident, fewer new NPP projects would have been cancelled. Our simulations show that an immediate, costless extension in operating licenses to 60 years would double the expected present discounted value of profits to continued operation of US NPPs.

---

[2] An example is the Shoreham NPP that was closed two months after receiving its operating license in 1989 because of the prospect of an extended shutdown to overcome opposition by New York State regulators to its safety evacuation plan.

There are several key differences between our analysis and previous economic analyses of the NPP decommissioning and life extension options. First, our analysis recognizes there is no "free disposal" in the nuclear power industry: NPP decommissioning is an expensive undertaking; see discussion in Pasqualetti and Rothwell (1991). It might be optimal to continue to operate a higher cost NPP than to immediately replace it with a lower cost fossil fueled plant. Several previous studies did not explicitly account for the high costs of NPP decommissioning, which imply that an uneconomic radioactive NPP has a large negative private and social value.

Second, our analysis recognizes that the historical construction costs for existing NPPs are *sunk costs* and thus irrelevant to the decision of whether to continue to operate a plant in the future.[3] Many comparisons (e.g., Hewlett, 1991a) include the "opportunity cost of capital" as part of NPP "revenue requirements" in their calculations of the costs and benefits of NPP life extension. However, this treatment presumes that an uneconomic NPP is a piece of capital that can be costlessly transformed into other productive uses. We argue that cost/benefit calculations that include the annuitized value of sunk NPP construction costs are improperly biased toward fossil fuel alternatives. For example, in 1990 the total generation expenses for the median cost NPP were 4.69 cents per kilowatt hour compared to 2.76 cents per kilowatt hour for a median cost coal plant (EIA, 1992). The cost of 4.69 cents for nuclear power includes 2.79 cents per kilowatt hour of capital related expenses, i.e., the annuitized sunk investment costs, taxes, depreciation, etc. If we restrict the comparison to *unsunk* operating costs (e.g., operating, maintenance, and fuel expenses), nuclear power is still cheaper than fossil fuels: 1.81 cents per kilowatt hour for the median cost NPP versus 2.25 cents per kilowatt hour for the median cost coal plant. We argue that only the unsunk operating and maintenance costs together with the expected discounted costs of any future capital additions necessary to keep the plant running safely are relevant for studying prospective operating decisions of *existing* NPPs.

A third difference is that our analysis treats the decision to decommission an existing NPP as logically distinct from the decision to invest in new generating capacity. Previous studies, such as Hewlett (1991b), determine the optimal time to close a plant as the point at which the present value of costs of operating a plant exceed the present value of costs of supplying equal power from the best alternative source. However, it is possible for a utility to continue to operate its NPP *and* invest in a new generating plant: the two alternatives need not be mutually exclusive. Underlying these previous calculations is the presumption that utilities face perfectly inelastic demands for electricity from their local service areas and do not have access to power grids to sell excess generating capacity. In this case it does not make sense to have any excess generating capacity beyond what is necessary to meet peak local demand. However, our analysis presumes that utilities can freely buy or sell power over power grids at a competitive regional price of electricity. In such a world the price of electricity provides the correct shadow price governing investment, operating and closure decisions. This enables us to treat investment and operating decisions of each generating plant

[3] This is not quite true. Historical sunk costs might be relevant when state public utility commissions allow the operator to recoup undepreciated NPP construction costs in the form of higher future electricity prices. Our estimated DP model accounts for this possibility.

separately. In particular it can be economic for a utility to continue to operate a high cost NPP and simultaneously invest in a new lower cost fossil fuel plant though the combined output of both plants exceeds peak local capacity: the utility simply sells its excess power to other utilities over the power grid. In a world of competitive power pooling each generating station becomes an independent "profit center" and it will be optimal to operate such a plant until the time when the present value of profits (or losses) from continued operation exceeds the present value of the costs of closing the plant. The two approaches, isolated local electric markets versus our assumption of competitive regional markets for electricity are idealizations, but we feel the latter is a better approximation to the present and future structure of the US electrical power industry.

A final difference is that our estimates of the potential gains to extended license spans are not based on direct measurements of electricity revenues and the fuel, operating, and maintenance costs incurred by NPPs. Instead we adopt a "revealed preference" approach in which we infer NPP profit functions by maximum likelihood estimation of the unknown parameters of an unrestricted specification of a DP model of optimal operation of a plant using monthly observations of NPP operations. The maximum likelihood parameter estimates yield an estimated DP model that closely mimics operating characteristics of actual NPPs. While the maximum likelihood approach is unable to identify the location and scale of NPP profits, by re-solving the DP model under the hypothesis of 60 license lifetimes and computing the ratio of the estimated value functions we are able to identify the *relative* increase in expected discounted profits resulting from a license extension. In addition the DP model yields detailed predictions of optimal operating decisions under the new policy regime. Previous studies did not employ a formal econometric model of NPP operations, but based their estimates on accounting measurements of NPP fuel, operating, and maintenance costs. Each approach has strengths and weaknesses. A potential weakness of our approach is that our estimates of the gains to license extension depend on a particular model of NPP operations, so our estimates could be invalid if our specification is incorrect. A weakness of the alternative approach is that accounting numbers might not accurately measure the true economic costs to running a plant. For example, most estimates of NPP operating costs ignore overhead costs like nuclear liability insurance and staff benefits. Some studies have estimated these omitted costs to be as high as 30% of total operating and maintenance costs (Office of Technology Assessment, 1993, p. 92). In future work we intend to integrate the two approaches to improve the accuracy of our estimates and provide additional tests of the validity of our structural econometric model.

Section 2 provides a brief overview of nuclear power generation and regulation, and discusses some of the age-related deterioration problems experienced by NPPs, and the uncertainties and costs 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 5. Section 3 presents our dynamic programming (DP) model of optimal operation of a plant. 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 the Graybook data and some of which are unobserved by the econometrician. Our DP model is designed to accommodate both types of signals. In Section 4 we use monthly data on US NPPs 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 5 presents simulation results from our estimated DP model, including predictions of industry output under the two different licensing scenarios summarized above. The last section presents concluding remarks and directions for future research.

2. Nuclear Power Technology, Plant Aging, and Prospects for License Extension

There are many types of NPPs in the world, but in the US nearly all commercial NPPs use either pressurized water reactors (PWRs) or boiling water reactors (BWRs). Of the 111 licensed US NPPs operating in 1993, 76 were PWRs (with reactors built by Babcock & Wilcox, Combustion Engineering, and Westinghouse) and 35 were BWRs (built by General Electric). PWRs and BWRs are types of light water reactors (LWRs), using ordinary water as coolant and moderator. LWRs generate power via nuclear fission using slightly enriched uranium as fuel. 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 NPP 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 tradeoff between (1) the potential improvement in capacity factor associated with longer operating cycles and (2) the potential increased risk of unplanned mid-cycle outages due to fuel and other failures. We consider how this tradeoff changes as NPPs age. NPPs experience two sorts of aging problems: (1) short-term aging (for example, fuel rod failures that increase with the duration of operating spells) and (2) long-term aging problems with the reactor vessel, steam generators (in PWRs), and with the NPP's cooling, instrumentation, and control systems. Refueling outages partially regenerate the within-cycle deterioration associated with burnup of nuclear fuel. It is more difficult to quantify the impact of long-term aging problems because of plant-specific learning-by-doing effects (see Lester and McCabe 1993) and general technological improvements in fuel reliability, instrumentation, and other aspects of nuclear power technology. (On the influence of organizational structure, see Rothwell, 1996.) The rate and duration of unplanned outages decrease monotonically with NPP age (see, for example, Rothwell and Rust 1995). Part of the problem with capturing aging effects is that much of the age-related degradation in an NPP can occur toward the end of its operating life, but there are few observations on NPPs more than 20 years old. Another problem is that while the risk of some failures is increasing with age, safety procedures are designed to make failures rare events that make it difficult to estimate their hazard rates from a few observations.

There is a large engineering literature on problems and operating strategies associated with age and utilization-related deterioration in NPPs (see, for example, Shah and MacDonald, 1992). Through direct measurements, this literature has identified key aging mechanisms and developed age-management strategies for dealing with them. Shah and MacDonald's ranking of the primary degradation sites and mechanism of the major LWR components identified damage to the reactor vessel from radiation embrittlement and boric acid corrosion as the 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 pressure 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 (as at Yankee Rowe). Although some studies have claimed that embrittlement problems can be reversed by thermal annealing and can extend the life of the reactor vessel by as much as 20 years (Dragonajtys, Griesbach, and Server, 1991), there has been little practical experience with this procedure and there is scientific uncertainty about the rate of re-embrittlement after annealing (Shah and MacDonald, 1992, 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. Age related deterioration in other NPP components has forced utilities to incur huge expenses to keep their plants operating. For example, corrosion problems have required replacement of steam generators in PWRs and of recirculation piping and pumps in BWRs. The cost of repairing or replacing these items is high and involves long downtimes. Having incurred these expenses, operators are more cognizant of the value of reducing the rate of deterioration by operating their plants conservatively. In addition to embrittlement, the reactor vessel suffers stress from temperature and pressure variations at startup and shutdown and during power transients. A NPP is designed to withstand a maximum number of transients over its lifetime.[4] Because of the stresses involved and high cost of downtime the operator will shut down a reactor only when it is absolutely necessary. However due to the many safety monitoring devices attached to NPPs, unplanned shutdowns are inevitable.[5]

---

[4] For example, Table 3.2 in Shah and MacDonald (1992) shows that over its 40-year life a PWR is designed to withstand 400 scrams from full power, 80 loss of flow or abnormal loss of load events, 2,000 step load increases of 10% of full power, and 15,000 power loading or unloadings at a rate of 5% of full power per minute.

[5] Many unplanned shutdowns experienced by NPPs are false alarms caused by failure of electronic monitoring devices and computer software. For example, Iowa's Duane Arnold unit went offline for 86 hours in June 1989 because a hand-held radio interfered with instrumentation. Rahn *et al* (1984, Table 12.16) shows that 28% of forced outages are a result of signals from engineered safety features. Problems in the instrumentation and control systems and in the electric power system account for another 22% of forced outages.

There is a presumption that the conditional probability of within-cycle failures (where the failure probability is a function of duration since last refueling) and aging-related failures (where the failure probability is a function of the age of the NPP) are *bathtub shaped*, i.e., failure probabilities are initially high, decrease to a minimum level, and then increase again. For example, the interaction of learning and gradual age-related deterioration can produce a bathtub-shaped pattern for age-related failures. However as discussed, it is difficult to detect the eventual upturn in the probability of age-related failures. Using data from the NRC's Graybooks (NUREG-0020), checked against the International Atomic Energy Agency's *Operating Experience with Nuclear Power Stations in Member States* (various years), Rust and Rothwell (1996) confirmed the existence of a bathtub-shaped pattern for within-cycle failures. In particular, there are significantly higher rates of unplanned outages after a cold startup in the first month following a refueling outage. This effect has been documented in statistical analyses by Stoller (1987, 1989).

It is even more difficult to identify the independent effect of age on plant reliability and O&M costs. Previous analyses (for example, Rothwell and Rust, 1995, and Energy Information Administration, 1995) have found systematic improvement in reliability and reductions in O&M costs with plant age, an effect that is usually interpreted as a combination of technological progress and learning-by-doing. Although we cannot confirm or reject the learning-by-doing hypothesis without a more complete investigation, we do not find strong evidence for a bathtub-shaped probability of forced outages as a function of NPP age. Figure 2.1 plots the monthly probability of forced outages as a function of plant age. We see sharp drops in the forced outage rate after the first several years of operation followed by more gradual decreases over the next 20 or so years of operation. The figure shows sharp upturns in the forced outage rate after 24 years and then sharp decreases after 28 years, but these patterns are likely to be artifacts of having so few observations on NPPs more than 24 years old. Indeed, in our sample there are only 4 NPPs that were more than 25 years old. Thus one needs to interpret the estimates of the tails of these age-related probabilities with caution.

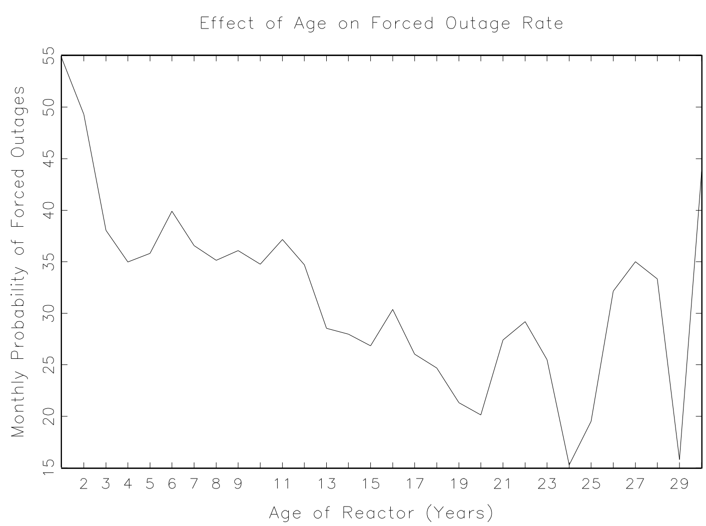

Figure 2.1 Effect of Reactor Age on the Probability of Forced Outages

We conclude this section with a discussion of the impact of Nuclear Regulatory Commission (NRC) safety regulations on NPP operating and decommissioning decisions. The NRC has the strongest influence of any regulatory body on the day-to-day operations of US NPPs. Under the Atomic Energy Act, the NRC is allowed to issue operating licenses for a maximum term of 40 years. Once a 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 is performed to protect the public and the environment from exposure to radioactivity. There are three strategies for reactor decommissioning: 1) immediate dismantlemant of the plant and disposal of radioactive waste and structures in long-term waste repositories; 2) initial decontamination followed by extended plant closure to allow radioactivity levels to decay to lower levels to permit safer and cheaper subsequent dismantlement; and 3) entombing the plant in a protective "sarcophagous" for up to 60 years followed by eventual release of the site. A current problem with the first decommissioning alternative is that due to local opposition and environmental challenges, the Department of Energy has not been able to secure a permanent repository for high level nuclear waste. As a result, estimates of the expected costs of decommissioning have increased substantially in the last decade. To date, the US has only had experience in decommissioning relatively small, experimental NPPs. Current estimates of the costs of decommissioning a typical 1 gigawatt NPP are in the range of several hundred million dollars, with a very large standard deviation (see Pasqueletti and Rothwell, 1991 and Office of Technology Assessment, 1993). In this paper we do not explicitly model the specific choice of decommissioning alternative and assume that the least cost alternative is chosen. Further, rather than attempt to incorporate direct information on the cost of decommissioning, our "revealed preference" approach infers operators' expectations of these costs from observed operating and closure decisions.

In 1991 the NRC began to draft a set of procedures that would allow operators the opportunity to apply for an extension of their operating license 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 and where the burden is on the regulator to prove a plant should have its license revoked because it can no longer be operated safely. Since 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 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 might believe that license renewal is not a realistic possibility. As we will show in Section 5, these beliefs could lead operators to close their plants well before the end of their 40-year license span. In view of the potential consequences of the large regulatory burdens and uncertainties surrounding the NRC's current license renewal procedures the Office of Technology Assessment suggested that "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 the plant could not continue to be operated safely. The issue of the appropriate policy regarding operating license extensions can be usefully analyzed within the context of a dynamic programming model of NPP operations. We introduce such a model in the next section. In Section 5 we use the DP model to show how costless license extension could increase the profitability and thereby extend the lifetime of the nuclear power industry.

3. A Dynamic Programming Model of NPP Operations

In this section we present a dynamic programming (DP) model designed to capture the main features of NPP operations discussed above. 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 NPPs are owned by utilities who are typically modeled as monopolists subject to rate of return regulation, theoretical work by Che and Rothwell (1995) and empirical evidence presented in Rust and Rothwell (1996) suggests that profit maximization provides a reasonable description of utility's objectives in the post-TMI regulatory environment because of the sharp rise in the probability of operating cost disallowances in public utility commission rate making decisions, and the increasing prevalence of other incentive-based regulations. In addition falling costs of fossil fuels during the 1980's put strong pressure on utilities to operate their NPPs in a cost-minimizing manner, subject to the more stringent safety regulation imposed by the NRC that we explicitly account for in our model. In addition, as we noted in the introduction, we presume that utilities are able to freely buy and sell electricity over power grids so NPPs are not subject to additional generation constraints from low local demand for electricity or lack of reliable sources of replacement power. Therefore, the current price of electricity is the relevant shadow price governing NPP operating decisions, enabling us to treat each plant 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. Our model does account for seasonal fluctuations in power demand and its implications for operating and shutdown decisions of NPPs.

We build on the DP model of Rust and Rothwell (1996) which assumed that an NPP is in either one of two spells at any point in time, an *operating spell* or a *refueling spell.* The model allows the operator to choose the level of NPP availability during an operating spell (including complete NPP shutdowns to repair unplanned equipment failures or for planned preventive maintenance) and the timing of each refueling outage. Constraints imposed by NRC safety regulations were modeled via "exogenous refueling durations," i.e., the duration of a refueling spell is random variable beyond the control of the NPP operator. Rust and Rothwell showed that the increase in NRC-mandated surveillance and preventive maintenance activities during refueling outages significantly increased the duration of refueling outages after the TMI accident since these activities typically take far longer to accomplish than the refueling of the reactor *per se*. Further, operators have large *ex ante* uncertainties about refueling durations: more than 40% of all refueling spells last 4 weeks longer than *ex ante* expectations. Refueling durations are the most important factors limiting achievable availability factors. Rust and Rothwell showed that the nuclear power industry reacted optimally to the increased refueling durations by increasing the planned duration of operating spells from 12 months to 18 months. Given the high opportunity cost of downtime (over $0.5 million per day in replacement power costs for a 1 gigawatt reactor), the increase in refueling durations are a significant factor contributing to the reduction in profitability of NPPs in the post TMI era.

In describing the state of the NPP, it is important to distinguish between outages that succeed in diagnosing and repairing problems in a short time versus protracted outages that can last for many months. The two-spell model of Rust and Rothwell had difficulty accounting for protracted outages, since the model had no way of "explaining" why it would be rational for the operator to voluntarily run the NPP at 0% capacity for an extended period. In this paper we extend the DP model of Rust and Rothwell (1996) to a three-spell model that can account for protracted outages. 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 from 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. An example is the discovery of small cracks or radiation-induced embrittlement of the reactor vessel following routine surveillance during a refueling outage. The prospect of an extended shutdown to repair such problems lead to the ultimate closure of Yankee Rowe in February 1992. Another example of involuntary major problem spells are the extended administrative shutdowns imposed by the NRC to correct management-related safety problems, such as the multi-year shutdowns at the Brown's Ferry plants. 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, i.e., 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 increase substantially.

We follow Rust and Rothwell (1996) and formulate the DP problem in discrete time. NPP operating decisions are assumed to be made at the start of each month which is convenient for matching the predictions of the DP model to the monthly Graybook data. 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, for example, adjustment of control rods and concentrations of moderators in the reactor coolant, the specific repair and inspection activities during forced outages, and the fuel management strategy during refueling outages. Although day-to-day management of NPPs involves complex tradeoffs, given the high opportunity costs of NPP downtime, these complexities are secondary to the larger long-run issues of operating NPPs. Our DP model is designed to capture the most important longer run tradeoffs that NPP operators face, i.e., the duration of operating spells, the timings of 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 availability 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 the Graybook, or that are even feasible to record. Therefore, we assume that the state variable, s_t, can be partitioned into two components $s_t = (x_t, \epsilon_t)$, 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 the operator can take. Thus, we will interpret ϵ_t as a vector with the same number of elements as the possible values of the control variable, a_t. Since 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, and 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 $\pi(a, x, \epsilon)$. It has the additively separable representation

$$\pi(a, x, \epsilon) = \mu(a, x, \phi) + \epsilon(a), \tag{3.1}$$

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 $\lambda(x_{t+1}, \epsilon_{t+1}|x_t, \epsilon_t, a_t)$ and that the NPP operator chooses an optimal operating strategy $a_t = \alpha_t(x_t, \epsilon_t)$ that maximizes the NPP's expected net present value $V_0(x, \epsilon)$ given by

$$V_0(x, \epsilon) = \max_{(\alpha_0, \dots, \alpha_T)} E\left\{ \sum_{t=0}^{T} \beta^t \pi(a_t, x_t, \epsilon_t) \Big| x_0 = x, \epsilon_0 = \epsilon \right\}. \tag{3.2}$$

In many DP problems the horizon T is often not well defined. However for an NPP, 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. However, in Section 5 we consider the empirical case for other expectational hypotheses and in Section 6 we examine the predictions of DP model under the hypothesis that operating licenses can be costlessly extended from 40 to 60 years.

We now describe the (observed) state and control variables used in our 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 and the operator receives a signal that the refueling spell or major problem spell has ended and the choice set is $A_t(x_t) = \{1, 3, 4, 5, 6, 7, 8\}$.

Implicit in our treatment of major problem spells is the assumption that the NPP is fully refueled and refurbished by the end of the major problem spell. This is why we rule out choice 2, enter a new refueling spell, at the end of both a refueling spell or a major problem spell. We argue that our 6 cell discretization of the continuous availability decision during operating spells is more than adequate for modeling the range of observed availability levels. A histogram of monthly availability levels reveals that most of the time a plant is running either at 100% or 0% availability, and the small mass of intermediate availability rates is uniformly distributed between 0 and 100%. (See Figure 3.1 in Rothwell and Rust, 1995, p. 24.) This is even more true when we view load curves in continuous time: a 50% availability rate for one month typically corresponds to running at 100% for half the month and 0% for the other half.

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, i.e., 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 $(f_{t+1}, \epsilon_{t+1})$ are realized, and the NPP operator makes the next decision in period $t+1$.[6]

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:

$$r_{t+1} = \begin{cases} 1 & \text{if } f_t = 3 \text{ and } (r_t = 3 \text{ or } r_t = 1) \\ 2 & \text{if } (a_t = 2 \text{ and } f_t < 3 \text{ and } r_t = 3) \text{ or } (r_t = 2 \text{ and } f_t = 3) \\ 3 & \text{if } a_t > 2 \text{ and } f_t < 3, \end{cases} \tag{3.3}$$

$$d_{t+1} = \begin{cases} d_t + 1\{a_t \neq 2 \text{ and } r_t \neq 3\} & \text{if } r_{t+1} = r_t \\ 1 & otherwise. \end{cases} \tag{3.4}$$

The law of motion for duration corresponds to "partial regeneration" of the NPP following each refueling or major problem spell. In the model we present below, rates of forced outages and operating costs increase with the

[6] 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 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 to focus attention on the more important issues of the output levels and timing of refuelings for which a monthly interval is appropriate. We believe that the errors arising 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 $\{\epsilon_t\}$ is *IID*). In future work we plan to adopt the continuous-time semi-Markov control framework of An (1993). This will enable us to avoid the measurement and interpretation problems arising from our discrete time approximation to the true continuous time control process.

duration of the operating spell. However, recalling the previous discussion of the "bathtub shaped" hazard function, rates of forced outages are high immediately after a refueling from problems associated with the reassembly of the reactor head, etc. 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. Additionally, Rust and Rothwell (1996) present 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 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. However because of technologically progress and learning-by-doing, the net effect aging is estimated to be negative, i.e., rates of forced outages decrease with the age. Mid-cycle preventive maintenance outages also can help to regenerate the NPP. This is reflected in the formula for duration, Equation (3.4): $d_t + 1$ is incremented for each period the NPP continues in the current spell $(r_{t+1} = r_t)$ except when there is a shutdown during an operating spell $(a_t = 2$ 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 net discounted costs involved in the NPP closure 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, i.e., $A_{480}(x) = \{1\}$. This assumption is relaxed in Sections 4 and 5.

The law of motion for the NPP status variable f_t is probabilistic. Its probability distribution is derived from 5 conditional probabilities:

p_{of} is the probability of one or more forced outages occurring during an operating spell;

p_{rf} is the probability of one or more forced outages occurring in the first month following a refueling outage;

p_{om} is the probability of entering a major problem spell from an operating spell;

p_{mo} is the probability of resuming operation from a major is the problem spell;

p_{ro} is the probability of resuming operation from a refueling outage.

Each of these conditional probabilities depend 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

$$p_i(x_t, a_t, t) = \frac{\exp\{g(x_t, a_t, t, \psi_i)\}}{1 + \exp\{g(x_t, a_t, t, \psi_i)\}}, \quad i = of, rf, om, mo, ro, \tag{3.5}$$

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 3 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

$$f_{t+1} = \begin{cases} 1 & \text{with probability } (1 - p_{om})(1 - p_{of}) \\ 2 & \text{with probability } (1 - p_{om})p_{of} \\ 3 & \text{with probability } p_{om}. \end{cases} \tag{3.6}$$

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

$$f_{t+1} = \begin{cases} 1 & \text{with probability } p_{mo}(1 - p_{of}) \\ 2 & \text{with probability } p_{mo}p_{of} \\ 3 & \text{with probability } (1 - p_{mo}). \end{cases} \tag{3.7}$$

If the operator initiates a refueling outage (i.e., if $a_t = 2$) or if the current month is a continuation of refueling (denoted by $r_t = 2$ and $f_t = 3$, as explained below), there is a similar law of motion for f_{t+1} as in Equation (3.7) but with p_{ro} replacing p_{mo} and p_{rf} replacing p_{of}.

The NPP's profit function $\pi(a_t, x_t, \epsilon_t)$ was specified in Equation (3.1). Let $u(a)$ denote the level of electricity generated by the NPP, given the availability decision, a. Let p_t denote the price of electricity at time t. The specification of the $\mu(a_t, x_t, \phi)$ containing the observed state variables is given by

$$\mu(a_t, x_t, \phi) = \begin{cases} -\phi_c & \text{if } a_t = 1 \text{ (close the NPP)} \\ -c_r(x_t, \phi_r) & \text{if } a_t = 2 \text{ (refuel the NPP)} \\ p_t u(a_t) - c_o(x_t, a_t, \phi_o) & \text{if } a_t > 2 \text{ (operate the NPP at level } a), \end{cases} \tag{3.8}$$

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.

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 (1996) and normalize the present value of decommissioning costs as 0, i.e., $\phi_c = 0$. To simplify our model, we assume that the price of electricity is constant, 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 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. While 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 this assumption in future work.

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

$$\lambda(x_{t+1}, \epsilon_{t+1} | x_t, \epsilon_t, a_t) = p(x_{t+1} | x_t, a_t, \psi) q(\epsilon_{t+1}), \tag{3.9}$$

where ψ is a vector of unknown parameters characterizing the transition density for the observable part of the state and control variables. Equation (3.9) is known as a "conditional independence assumption" since 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) showed that the conditional choice probabilities, $P_t(a|x)$, are given by the classical multinomial logit formula:

$$\begin{aligned} P_t(a|x) &= \int I\{a = \alpha_t(x, \epsilon)\} q(d\epsilon) \\ &= \frac{\exp\{v_t(x, a)\}}{\sum_{a' \in A_t(x)} \exp\{v_t(x, a')\}} \end{aligned}, \tag{3.10}$$

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

$$v_t(x, a) = \mu(x, a, \phi) + \beta \int \log \left[\sum_{a' \in A_t(x')} \exp\{v_{t+1}(x', a')\} \right] p(dx' \Big| x, a, \psi). \tag{3.11}$$

The v_t functions are related to the value function $V_t(x, \epsilon)$ given in (3.10) by the identity:

$$V_t(x, \epsilon) = \max_{a \in A_t(x)} [v_t(x, a) + \epsilon(a),] \tag{3.12}$$

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 using standard backward induction using the recursion Equation (3.11). The implied stochastic process for the observed state and control variables $\{x_t, a_t\}$ constitutes the DP model's prediction of the optimal operation 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 provides observations on the realization of the observed state and control variables for the sample of US NPPs that were operational at the start of our sample in January 1989. Denote this data by $\{x_t^i, a_t^i\}$, $t = \underline{t}_i, \dots, \overline{t}_i$, $i = 1, \dots, N$. Given the conditional choice probability $P_t(a|x)$ in Equation (3.10) and the decomposition of the transition density $\lambda(x', \epsilon' | x, \epsilon, a)$ in Equation

(3.9), it is straightforward to estimate the unknown parameter vector $\theta = (\beta, \phi, \psi)$ by maximum likelihood using the (full) likelihood function

$$L_f(\theta) = \sum_{i=1}^{N} \sum_{t=\underline{t}_i+1}^{\overline{t}_i} \log \left[P(a_t^i | x_t^i, \theta) p(x_t^i | x_{t-1}^i, a_{t-1}^i, \psi) \right]. \quad (3.13)$$

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

$$L_1(\psi) = \sum_{i=1}^{N} \sum_{t=\underline{t}_i+1}^{\overline{t}_i} \log \left[p(x_t^i | x_{t-1}^i, a_{t-1}^i, \psi) \right], \quad (3.14)$$

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

$$L_2(\beta, \phi | \hat{\psi}) = \sum_{i=1}^{N} \sum_{t=\underline{t}_i+1}^{\overline{t}_i} \log \left[P(a_t^i | x_t^i, \beta, \phi, \hat{\psi}) \right]. \quad (3.15)$$

Rust (1988) established the consistency and asymptotic normality of the two-stage and full maximum likelihood estimators.[7]

4. Estimation and Testing of the DP Model

This section present 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 the previous section. 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 4.1. 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 the previous section. Note that the unrestricted specification also includes monthly dummies to reflect seasonal variations in the opportunity cost of an outage. We also imposed an identifying normalization that the dummy coefficient for November is zero.[8] The first column of Table 4.2 presents the definitions of the individual components of the ψ coefficients (the unknown coefficients characterizing the law of motion $p(x_{t+1} | x_t, a_t, \psi)$ as described in the previous section).

---

[7] The covariance matrix for the parameters (β, ϕ) will not be consistently estimated from the second stage partial 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.

[8] We were able to conveniently include monthly dummies without a 12-fold expansion of the state space by synchronizing the plant age modulo 12. To accomplish this we made slight adjustments to the ages of each NPP to guarantee that whenever the plant age in months was evenly divisible by 12 the month in question would be December of the given calendar year. Thus, *month* $=$ *age mod* 12.

Parameter	Description
β	monthly discount factor (fixed at $\beta = .999$)
$\phi_{a=1}$	expected present discounted value of costs of decommissioning NPP (normalized to 0)
$\phi_{a=2}$	expected cost of refueling the NPP, $\mu(2, r, d, f, \phi) = \phi_{a=2}$
$\phi_{a=2,f=2}$	extra refueling cost following a forced outage, $\mu(2, r, d, 2, \phi) = \phi_{a=2} + \phi_{a=2,f=2}$
$\phi_{a=3,f=3}$	per period cost of NPP shutdown during a major problem spell, $\mu(3, r, d, 3, \phi) = \phi_{a=3,f=3}$
$\phi_{a=3,f<3}$	cost of NPP shutdown during an operating spell, $\mu(3, r, d, 3, \phi) = \phi_{a=3,f<3}$
$\phi_{a=3,f=2}$	extra cost of shutdown due to forced outage, $\mu(3, r, d, 1, \phi) = \phi_{a=3,f<3} + \phi_{a=3,f=2}$
$\phi_{d,u>0}$	effect of operating cycle duration on expected profits (given positive availability)
$\phi_{u\in(0,.25]}$	expected profit of availability between 0 and 25%, $\mu(4, 3, d, f, \phi) = \phi_{u\in(0,.25]} + d\phi_{d,u>0}$
$\phi_{u\in(.25,.50]}$	expected profit of availability between 25 and 50%, $\mu(5, 3, d, f, \phi) = \phi_{u\in(.25,.50]} + d\phi_{d,u>0}$
$\phi_{u\in(.50,.75]}$	expected profit of availability between 50 and 75%, $\mu(6, 3, d, f, \phi) = \phi_{u\in(.50,.75]} + d\phi_{d,u>0}$
$\phi_{u\in(.75,1)}$	expected profit of availability between 76 and 99%, $\mu(7, 3, d, f, \phi) = \phi_{u\in(.75,1)} + d\phi_{d,u>0}$
$\phi_{u=1}$	expected profit of 100% availability, $\mu(8, 3, d, f, \phi) = \phi_{u=1} + d\phi_{d,u>0}$
$\phi_{u=1,f=2}$	profit of 100% availability following forced outage signal, $\mu(8, 3, d, 1, \phi) = \phi_{u=1,f=2} + \phi_{u=1}$
$\phi_{dec}, \ldots$	adjustment to profit for an outage in December, January, etc.
$\psi_{rf}(1)$	constant term in probability of forced outage following refueling spell
$\psi_{rf}(t)$	age term in probability of forced outage following refueling spell
$\psi_{ro}(d_t = 1)$	dummy on first month in refueling spell duration model
$\psi_{ro}(d_t = 2)$	dummy on second month in refueling spell duration model
$\psi_{ro}(d_t = 3)$	dummy on third month in refueling spell duration model
$\psi_{ro}(d_t = 4)$	dummy on fourth month in refueling spell duration model
$\psi_{ro}(d_t \geq 5)$	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 4.1 Definitions of ϕ and ψ parameters

Following Rust and Rothwell (1996) we estimated the DP model using observations on a plant's *availability factor* (the fraction of the month the NPP is operating) rather than *capacity factor* (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). Availability factor is greater than capacity factor when a plant is operating at less than full utilization or is "load following," i.e., reducing power output to meet local demand constraints. Since most NPPs 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 re-estimate the DP model using observed capacity factors.[9]

Unlike Rust and Rothwell (1996) who estimated the DP model using a "post-TMI" sample covering the period January 1984 to December 1993, we estimate the DP model for a sample from January 1989 to December 1994, using operating data for 1994 that recently became available. Rust and Rothwell defined their "post-TMI" sample as beginning after 1984 rather than immediately following the TMI accident in March 1979 to account for the fact that it took time for regulatory policy to adjust and stabilize following the accident. In particular, it took over one year for the Kemeny Commission to conduct an examination of the causes of the TMI accident, and even longer for the recommendations of this report (which included much more stringent oversight by the NRC) to be fully implemented and finalized. Rust and Rothwell provided evidence that the industry was in a state of flux from 1979 to 1984. Although NRC and state PUC regulatory policies had solidified by 1984, NPP operating and maintenance (O&M) costs continued to rise 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 (Energy Information Administration, 1995, p. 29, and Office of Technology Assessment, 1993, p. 24). However since 1989, real O&M costs have stabilized, rising at an annual real rate of less than 1%. As noted in the introduction, calculations by Hewlett (1991a) and others showed that if both the levels and rate of growth of O&M costs that prevailed during the mid-1980s continued, it would uneconomic to extend the life of most NPPs beyond 40 years, and indeed it would be optimal to close many higher cost NPPs before the end of their license expiration dates. Therefore, to provide the most optimistic case for nuclear power, 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.

Tables 4.2 and 4.3 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 (3.13). The estimates are conditioned on a fixed value of the discount factor, $\beta = .999$, that corresponds to a real annual interest rate of 1.2%. As in previous work, it is difficult to identify the exact value of β precisely, a likelihood ratio test decisively rejects the hypothesis that operators use discount rates higher than 10% per year. The results tell us that operators

[9] See Rothwell 1990 for more detailed discussion of the capacity factor.

Parameter	Estimate	Standard Error	t-statistic
$\phi_{a=2}$	-2.636	0.341	-7.72
$\phi_{a=2,f=2}$	-3.704	0.295	-12.57
$\phi_{a=3,f=3}$	-1.768	0.667	-2.65
$\phi_{a=3,f<3}$	-2.220	0.194	-11.46
$\phi_{a=3,f=2}$	-4.408	0.331	-13.31
$\phi_{d,u>0}$	$-.703 \times 10^{-1}$	$.145 \times 10^{-1}$	-4.84
$\phi_{u\in(0,.25]}$	-3.527	0.199	-17.74
$\phi_{u\in(.25,.5]}$	-3.104	0.188	-16.51
$\phi_{u\in(.5,.75]}$	-2.194	0.177	-12.40
$\phi_{u\in(.75,1)}$	-1.047	$.171$	-6.11
$\phi_{u=1}$	1.548	0.168	9.24
$\phi_{u=1,f=2}$	-5.934	$.195$	-30.41
ϕ_{dec}	$-.692$	0.291	-2.38
ϕ_{jan}	$-.636$	0.260	-2.45
ϕ_{feb}	$-.743$	0.231	-3.21
ϕ_{mar}	0.353	0.223	1.58
ϕ_{apr}	0.336	0.237	1.42
ϕ_{may}	0.143	0.239	0.60
ϕ_{june}	$-.521$	0.254	-2.05
ϕ_{july}	$-.694$	0.255	-2.72
ϕ_{aug}	-1.551	0.242	-6.41
ϕ_{sep}	0.210	0.243	0.87
ϕ_{oct}	0.241	0.288	0.84

$\log(L_f(\hat{\phi}, \hat{\psi})) = -9011.66$

$N = 7526$

Table 4.2 Full information maximum likelihood estimates of parameters of ϕ parameters

seem to be extremely farsighted in their decision making, and the fact that the likelihood function is essentially flat for discount factors greater than $\beta = .999$ might be telling us that their objective is to maximize long-run average profits (see Rust, 1987, for a similar finding in the context of a monthly model of bus engine maintenance and replacement).

Because parameter estimates in Table 4.2 are similar to those in Rust and Rothwell (1996), we concentrate on describing the findings we have obtained for our new sample and our new 3-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. However, the relative values of the various coefficient estimates are meaningful in this model. Starting with the coefficient of the profit corresponding to 100% availability under "normal" conditions, $\hat{\phi}_{u=1} = 1.548$, we see that profits corresponding to successively lower levels of availability, $\hat{\phi}_{u\in(.75,1)}, \cdots, \hat{\phi}_{u\in(0,.25]}$

decrease monotonically, as expected. The only anomaly is that the profit (loss) corresponding to a complete shutdown of the NPP for the month, $\hat{\phi}_{a=3,f<3} = -2.22$ is lower (higher) than the profit (loss) corresponding to running the NPP at availability factors from 1 to 75%. This is similar to the finding in Rust and Rothwell (1996) and we believe the explanation for the anomaly is the same: since 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. This would explain the anomalous finding that a complete shutdown of the NPP is estimated to be more profitable 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 reflects the regeneration that occurs from preventive maintenance during a refueling outage is to reset the spell duration counter d_t back to 1. We acknowledge that our estimate of $\hat{\phi}_{a=2}$ might be biased upward if our 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 the next operating spell (i.e., 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).

However, we note that our estimation results reveal significant declines in profits in months where forced outage signals are received ($f = 2$). For example, the estimation results show that the profit (loss) of shutting down the NPP or initiating a refueling outage in response to a forced outage is significantly lower (higher) than the profit (loss) of running at any positive availability level. For example, the monthly loss resulting from initiating a refueling following a forced outage signal is $-6.3 = \hat{\phi}_{a=2} + \hat{\phi}_{a=2,f=2}$, and the loss required to shut down the NPP to repair damage resulting from a forced outage is $-6.6 = \hat{\phi}_{a=3,f<3} + \hat{\phi}_{a=3,f=2}$. Furthermore, our estimation results reveal that NPP operators are highly averse to "imprudent" operation of their NPPs, in the sense of insisting on running 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 lower (higher) than the profit (loss) of 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 midcycle preventive maintenance outage. This also could be an anomalous finding because expensive capital upgrades and retrofits typically occur during major problem spells (for example, replacement of steam generators). Once again, 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

Parameter	Estimate	Standard Error	t-statistic
$\psi_{rf}(1)$	$-.847$	0.223	-3.80
$\psi_{rf}(t)$	0.107×10^{-2}	0.129×10^{-2}	0.83
$\psi_{ro}(d_t = 1)$	-5.928	1.004	-5.91
$\psi_{ro}(d_t = 2)$	-1.750	0.145	-12.10
$\psi_{ro}(d_t = 3)$	0.132×10^{-1}	0.112	0.12
$\psi_{ro}(d_t = 4)$	0.474	0.162	2.93
$\psi_{ro}(d_t \geq 5)$	$-.137$	0.408	$-.33$
$\psi_{ro}(\quad (d_t - 4)(d_t \geq 5) \quad)$	0.267	0.226	1.18
$\psi_{of}(1)$	-1.444	0.106	-13.58
$\psi_{of}(t)$	$-.441 \times 10^{-3}$	0.352×10^{-3}	-1.25
$\psi_{of}(d_t)$	0.320×10^{-1}	0.222×10^{-1}	1.44
$\psi_{of}(d_t^2)$	$-.210 \times 10^{-2}$	0.124×10^{-2}	-1.70
$\psi_{of}(f_t = 1)$	0.340	0.755×10^{-1}	4.50
$\psi_{om}(1)$	-6.905	0.762	-9.06
$\psi_{om}(t)$	$-.548 \times 10^{-2}$	0.176×10^{-2}	-3.12
$\psi_{om}(d_t)$	0.166	0.392×10^{-1}	4.24
$\psi_{om}(f_t = 1)$	1.259	0.479	2.63
$\psi_{mo}(1)$	-3.350	0.556	-6.02
$\psi_{mo}(d_t)$	0.499×10^{-1}	0.307×10^{-1}	1.62

$\log(L_f(\hat{\phi}, \hat{\psi})) = -9011.66$

$N = 7526$

Table 4.3 Full information maximum likelihood estimates of parameters of ψ parameters

levels of O&M costs that would have been incurred if the capital upgrades and other repairs had not been undertaken. Furthermore, in the next section we present plots of the estimated value function that show the present value of profits (i.e., the value function V) is far lower during a major problem spell than during an operating or refueling spell). That is, the estimates imply that major problem spells entail a huge reduction in discounted operating profits.

In view of these problems one cannot interpret the individual profit function estimates in Table 4.2 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 since our interest focuses on recovering the value function V rather than the current profit function μ. We argue that our unrestricted specification of μ allows us to consistently estimate V though we have not fully specified the dynamic structure of how current maintenance and capital upgrade investments influence future O&M costs.

The remaining coefficients in Table 4.2 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 shutdown 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).

Table 4.3 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 Rust and Rothwell (1996) for the sample period 1984 to 1993, except (1) the estimates imply the refueling durations were significantly shorter on average after 1989 than they were from 1984 to 1988; (2) the "bathtub shaped" hazard function for forced outages has flattened out: rates of forced outages are almost constant at 20% per month, independent of the duration of the operating spell; (3) the level of serial correlation in forced outages has decreased; and (4) 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 2, it is difficult to detect aging effects in our sample of NPPs. Older NPPs appear to have a systematically lower risk of forced outages and major problems than younger NPPs. Apparently the combined effects of technological progress and learning-by-doing outweigh the effects of age-related deterioration.

We now evaluate the ability of the DP model to mimic NPP operations. Figure 4.1 presents a comparison of actual versus simulated operating histories for the Arkansas One, Unit 1 NPP from January 1984 to December 1990. The DP model is a "realistic" model of NPP operations in the sense that it is difficult to distinguish true observations of NPP operations from simulated observations generated on the computer — at least when the number of observations is small. However, our full dataset has 7426 observations and using these observations one can construct powerful statistical tests that are capable of rejecting the hypothesis that our DP model provides an "exact" model of NPP operations.

Table 4.4 presents the results of a Chi-square 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 6131 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:

$$\begin{aligned} \hat{P}(a|X) &= \int_{x \in X} \hat{P}(a|x)\hat{F}(dx|X) \equiv \frac{1}{N}\sum_{i=1}^{N} I\{a_i = a, x_i \in X\} \quad \text{(NP)} \\ P(a|X,\hat{\theta}) &= \int_{x \in X} P(a|x,\hat{\theta})\hat{F}(dx|X) \equiv \frac{1}{N}\sum_{i=1}^{N} P(a|x_i,\hat{\theta}) I\{x_i \in X\} \quad \text{(DP)} \end{aligned} \tag{4.1}$$

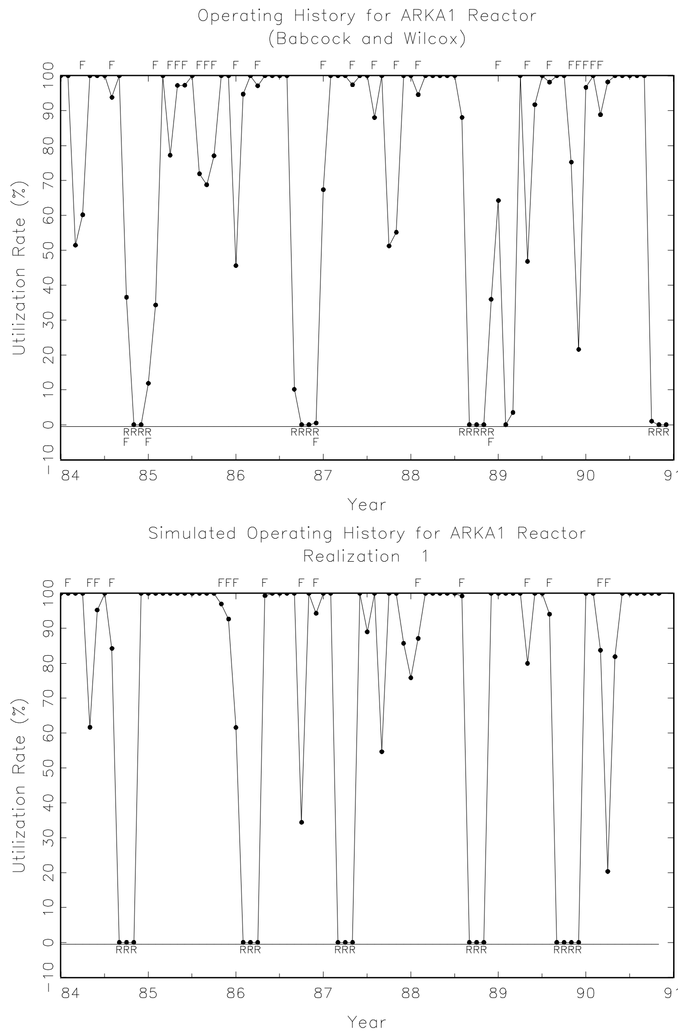

Figure 4.1 Actual and Simulated Operating Histories for Arkansas One, Unit 1 NPP

Andrews' (1988) version of the Chi-square 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 $\left(P(a_i|X_j, \hat{\theta}) - \hat{P}(a_i|X_j)\right)$ is to zero. The results in Table 4.4 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-square statistic has an asymptotic Chi-square distribution with 7 degrees of freedom.[10]

[10] Andrews' (1988) formulas for the covariance matrix of the residuals presumed that the true data generating process is *IID*. In our case if our specification is correct the data generating process is Markovian and not *IID*. However, it is straightforward to verify using the conditional independence properties of Markov processes that Andrews' formulas for the covariance matrix of the test statistic are valid when the data generating mechanism is Markovian.

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

$\chi^2(7) = 53.19$

Table 4.4 Predicted versus Actual Choice Probabilities: Full Sample

While the predicted and actual choice probabilities are quite close to each other, the magnitude of the Chi-square statistic indicates a decisive rejection of the DP model.[11] Reasons the DP model is rejected include (1) NPP operators might not be discounted profit maximizers; (2) there might be plant level heterogeneity that is not accounted for by the DP model; and (3) the conditional independence assumption Equation (3.9) might not be valid. Unfortunately the Chi-square statistic by itself is not informative as to which of the assumptions underlying the DP model are most likely to be violated.

Although rejected under this test, our DP model is able to mimic aggregate behavior of the nuclear industry. [12] Figure 4.2 shows this. We ran 20 stochastic simulations of the estimated DP model using the state of the nuclear power industry in January 1984 as initial condition.[13] We refer to these as "out-of-sample" simulations since 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 this is a good test of the ability of the DP model to make long-run forecasts since the divergence between predicted and actual trajectories usually increases with time in complicated nonlinear stochastic models.

---

[11] We experimented with alternative partitions $(X_1, \ldots, X_J)$ and found the results of the Chi-square test to be sensitive to the choice of the partition, the number of partition elements J and the choice of alternative asymptotically equivalent versions of the Chi-square statistic. We can find partitions where the DP model is not rejected. Therefore, we do not find the magnitude of the Chi-square statistic *per se* to be extremely informative.

[12] We refer the reader to Rust and Rothwell (1996) for more detailed comparisons of predicted versus actual behavior at the level of individual NPPs. One of the key findings from that paper is that the DP model correctly predicts the shift in planned durations of operating spells from 12 months before TMI to 18 months in the post-TMI period. The DP model "explains" this shift as a rational response to the increase in refueling durations due to the NRC's stricter safety regulations in the post-TMI period.

[13] If an NPP came online after January 1984 we began the simulation from the first month that it became operational.

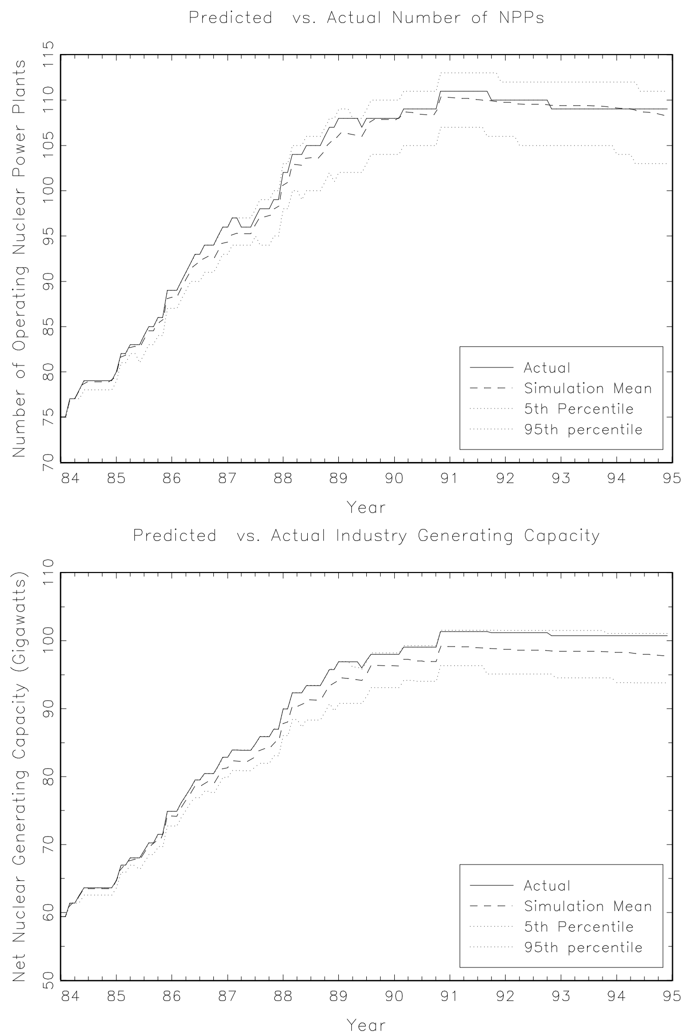

Figure 4.2 Comparison of Actual vs. Simulated Number of Plants and Total Generating Capacity

The top panel of Figure 4.2 plots the simulated versus actual number of operating NPPs. The mean number of NPPs in the 20 simulations tracks the actual number of NPPs well. Also, we plotted 90% confidence bands (computed pointwise as the highest and lowest number of firms in the 20 simulations at each point in time). Because these bands are tight, predictions of a rash of "early retirements" seem unfounded. The bottom pane of Figure 4.2 plots simulated versus actual industry generating capacity. In this case the mean capacity in the 20 simulations slightly underpredicts capacity levels and the 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 4.2 is that the 6 NPPs that were closed since 1986 were smaller than the average NPP.[14]

[14] The average generating capacity in our sample is 865 megawatts. The LaCrosse, Fort St. Vrain, and San Onofre 1 NPPs were 50, 330, and 410 megawatts, respectively.

5. Predicting the Lifetime of the Nuclear Industry Under Alternative Licensing Rules

In this section we use the estimated DP model to predict the evolution and decline 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 that 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 2, it would be necessary to model explicitly the decision of whether to apply for a license extension as part of the overall dynamic programming problem. The solution to this problem will depend on the operator's perceptions of the expected costs of submitting an application and the likelihood 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 about the regulatory and economic environment, including: (1) the NRC and PUC regulatory regimes will not change;[15] (2) the real price of electricity will remain at current levels; (3) the government will not impose carbon taxes or costly environmental regulations on fossil fuels; (4) real NPP O&M costs will remain at current levels; (5) the Department of Energy will be successful in developing a long-term waste repository for spent nuclear fuel; (6) there are no major changes in the expected costs of NPP decommissioning; and (7) 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 5.1 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. The top panel of Figure 5.1 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 a total of 20 NPP closures will occur between now and 2010. Note, however, that we predict that only 3 NPPs will be closed between 1995 and 2000. This result contrasts predictions that as many as 25 NPPs will be closed in the next few years. By the time the last NPP closes in 2031, we predict a total of 75 NPPs will have closed before their license expiration dates, leaving only 32 NPPs to be operated for the full duration of their operating licenses. The second panel of Figure 5.1 plots our prediction of the path of total nuclear generating capacity. For comparison purposes we also plot a "naive prediction" that assumes all NPPs will be operated continuously until their operating licenses expire. We see that the naive prediction is significantly higher than the simulated prediction from the DP model, with the gap between the two predictions rising to 120% of simulated

[15] We do not model the deregulation of the electricity generating industry and the introduction of competition. This is left for future research.

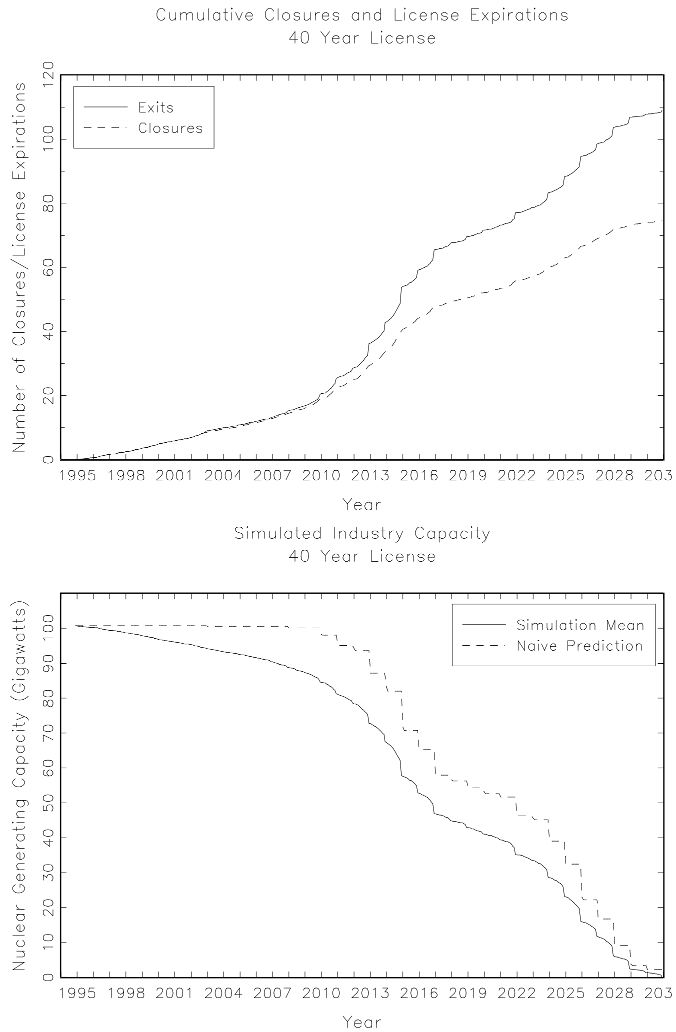

Figure 5.1 Simulated Evolution of the Nuclear Power Industry under 40 Year Operating Licenses

capacity by 2010 and reaching 147% of simulated capacity by 2021. Of course, the reason that the naive prediction is too high is that it takes no account of early retirements of NPPs.

Figure 5.2 presents forecasts of the evolution of the nuclear power industry under a 60-year license span. Figure 5.2 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. However it is more surprising to note 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 NPPs are closed more the 10 years "prematurely" under the regime of 40-year operating licenses. However, under 60-year operating licenses, only 3% of the NPPs close early.

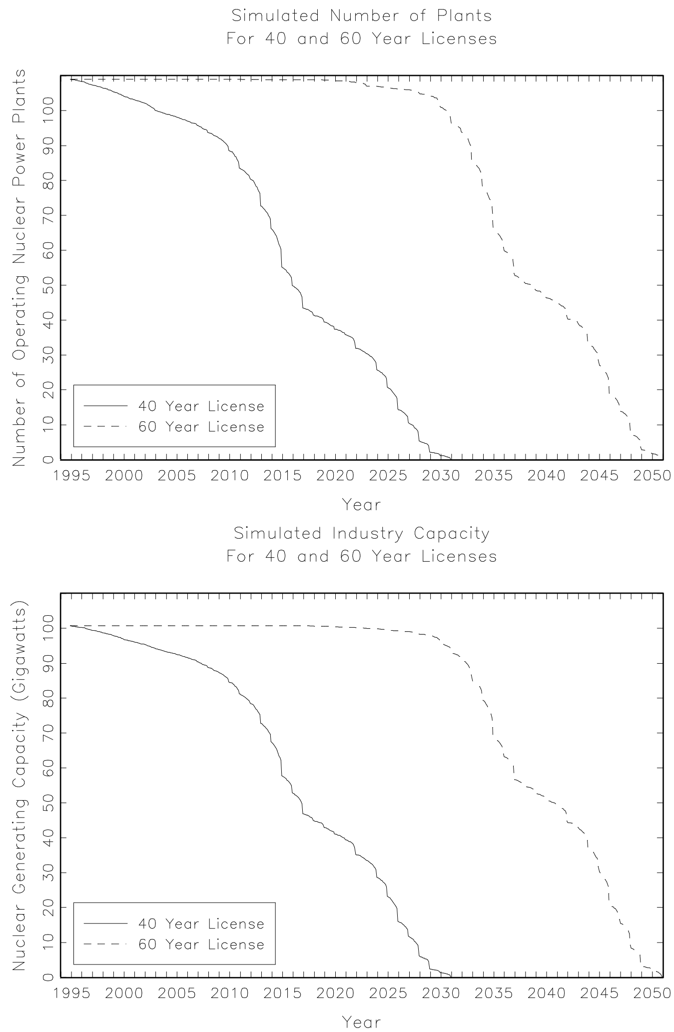

Figure 5.2 Predicted Evolution of the Nuclear Power Industry under 40 and 60 Year Operating Licenses

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 (i.e., to pay back the capital investments made to extend plant longevity). This is verified in Figure 5.3 that 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% 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 since there are no age terms entering the per period profit function (see the definition of the ϕ coefficients in Table 4.1). Because the probability of major problems and forced outages decreases with the plant age, NPPs become steadily more profitable per unit of time as they age. This fact is responsible for the initial upward slope in the value functions. However there is also a "horizon effect" that forces value functions to decrease after some point: as the NPP approaches its license expiration date the expected discounted profits fall. Thus, the quadratic shape of the value function is the result of the tradeoff between the increase in NPP profitability with age and the horizon effect.

Figure 5.3 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. However both panels show 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 increases substantially. Under a 40-year license the optimal policy is to close the plant if a major 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. However under a 60-year license, 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.

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 5.3 that moving to a 60-year license doubles the value function. Thus, we conclude that a 60 license span would double the (discounted) profits of the nuclear power industry.

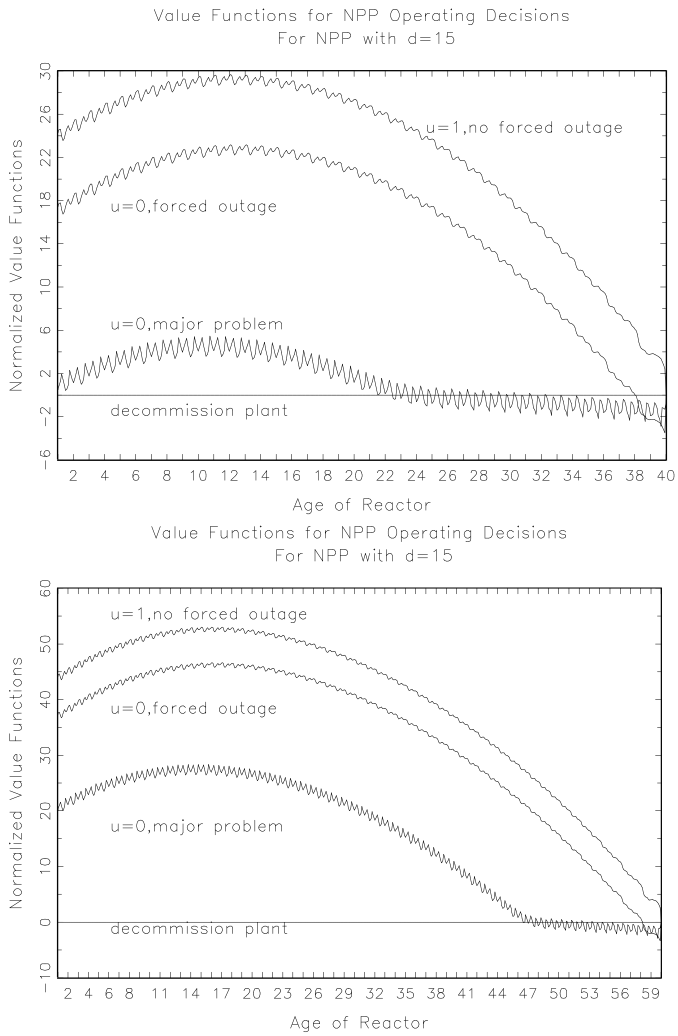

Figure 5.3 Value Functions under 40 and 60 Year Operating Licenses

6. Concluding Remarks

The policy simulations conducted in Section 5 are subject to three important qualifications. First, the accuracy of the predictions depend on the assumption 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 NPPs. Second, the predictions were based on the assumption there is no critical aging threshold, i.e., 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 may be much less than we have predicted. Third, the accuracy of our predictions depend on the accuracy of the DP model as a positive model of NPP operations. Specification tests performed in Section 4 indicate that the DP model is misspecified. One likely source of specification error is the homogeneity assumption: i.e., 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. However, if the differences in performance arise from suboptimal behavior by some NPP operators, a fundamental assumption underlying the DP model is called into question.[16]

The DP model presented in this paper is oversimplified in many respects. There are many directions in which it could be extended to better account for observed and unobserved heterogeneity. Further, this paper has been concerned with the question of determining the optimal lifetime of a NPP from the private perspective of the NPP owner rather than the social perspective of the regulator. 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. However, we believe our analysis provides a starting point for addressing this question. Given an appropriate social welfare function that represents society's preferences regarding the cost, safety, and environmental consequences of alternative forms of power generation, the design of optimal regulatory policy could be regarded as the solution to a Stackelberg game between utilities and regulators. We have taken the view that utilities should be regarded as behaving in a privately optimal manner, i.e., they can be modeled as adopting best responses to any given regulatory regime. Once we have estimated the "structural" parameters governing their preferences and production technologies, it is in principle possible to consider solving the Stackelberg game to characterize the optimal regulatory policy. However, we feel that this approach is overly ambitious until we provide a more satisfactory solution to the problem of identifying the policy invariant parameters of NPP operations and the NPP operators' preferences.

---

[16] Following Rust (1995), the hypothesis that utilities are behaving optimally is untestable without strong restrictions on their beliefs and objective function. Rust's results show that essentially any behavior pattern can be "rationalized" as optimal for some choice of parameters. However there is no guarantee that any choice of parameter values that does rationalize behavior will be "reasonable." To the extent that our estimates are regarded as reasonable, our results support the hypothesis that utilities are behaving optimally.

7. References

An, M. (1993) "Econometric Analysis of Sequential Discrete Choice Models," unpublished manuscript, Duke University.

Andrews, D.W.K. (1988) "Chi-Square Diagnostic Tests for Econometric Models: Theory" *Econometrica* **56-6** 1419–1454.

Campbell, J.L. (1988), *Collapse of an Industry: Nuclear Power and the Contradictions of U.S. Policy*, Ithaca, New York: Cornell University Press.

Che, Y.K. and G.S. Rothwell (1995) "Performance-Based Pricing at U.S. Electric Utilities, *The Energy Journal* **16-4**.

Cohen, B.L. (1990), *The Nuclear Energy Option*, New York: Plenum Press.

Dragonajtys, R.W., Griesbach, T.J., and Server, W.L. (1991), "Use of a Decision Tool for Assessing Nuclear Reactor Vessel Embrittlement Options," in *Service Experience in Operating Plants 1991* (PVP Vol. 221), eds. W.H. Bamford *et al*, New York: American Society of Mechanical Engineers.

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

Energy Information Administration (1992) *Electric Plant Costs and Production Expenses 1990* Washington, D.C.: U.S. Department of Energy (DOE/EIA-0455).

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

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

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

Hewlett, J.G. (1991b) "Financial Implications of Early Decommissioning" *The Energy Journal* **12** (Special Issue on Nuclear Decommissioning Economics) 279–291.

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

International Atomic Energy Agency (1974-1993), *Operating Experience with Nuclear Power Stations in Member States*, Vienna.

Lester, R.K., and M.J. McCabe (1993) "The Effect of Industrial Structure on Learning by Doing in Nuclear Power Plant Operation" *The RAND Journal of Economics Journal* **24, 3** 418–438.

Makovich, L. Forest, L. and T. Fletcher (1988) "U.S. National and Regional Impacts of Plant Life Extension" Sandia National Laboratories, SAND87-7136.

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

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

Rahn, F., Admamantiades, A., Kenton, J., and Braun, C. (1984), *A Guide to Nuclear Power Technology: A Resource for Decision Making*, New York: John Wiley.

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

Rothwell, G.S. (1996), "Organizational Structure and Expected Output for Nuclear Power Plants, *Review of Economics and Statistics* (forthcoming, latest draft November 1994).

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

Rust, J. (1987), "Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher," *Econometrica*, 55-5, 999-1033.

Rust, J. (1988), "Maximum Likelihood Estimation of Discrete Choice Processes," *SIAM Journal of Control and Optimization*, 26-5, 1006-1024.

Rust, J. (1995), "Structural Estimation of Markov Decision Processes" in *Handbook of Econometrics, Vol 4*, eds. R. Engle and D. McFadden, New York: North Holland.

Rust, J. and G.S. Rothwell (1996), "Optimal Response to Shift in Regulatory Regime: The Case of the U.S. Nuclear Power Industry," *Journal of Applied Econometrics*, (forthcoming, latest draft November 1995).

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

Stoller Corporation (1987), "The Influence of Fuel-Cycle Duration on Nuclear Unit Performance," Palo Alto, California: Electric Power Research Institute (NP-5042).

Stoller Corporation (1989), "The Influence of Fuel-Cycle Duration on Nuclear Unit Performance: An Update," Palo Alto, California: Electric Power Research Institute (NP-6333).

U.S. Nuclear Regulatory Commission (1989-1994), *Licensed Operating Reactors: Status Summary Report* (NUREG-0020).