A Dynamic Programming Model of U.S. Nuclear Power Plant Operations

Geoffrey Rothwell
Stanford University
`rothwell@leland.stanford.edu`

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

February 1995

Abstract: This paper presents a dynamic programming (DP) model of an electric utility's optimal policy for operating a nuclear power plant (NPP). The utility chooses the level of capacity utilization of the NPP as a function of signals about the NPP's current operating state. In each period the utility must determine whether or not to operate the reactor, or shut it down for preventive maintenance or refueling, or to permanently close the plant. Maintenance performed during periodic refueling outages partially "regenerates" the NPP, reducing the risk of unplanned forced outages in succeeding periods. However since NPP's are designed to satisfy base power load, refueling outages involve a substantial opportunity cost in terms of lost power generation. The utility faces a tradeoff between cost-efficiency, fuel-efficiency and plant safety and makes its decisions subject to exacting safety regulations by the U.S. Nuclear Regulatory Commission (NRC). Tightened safety regulations following the Three Mile Island accident in March 1979 have contributed to a near doubling of the mean duration of refueling outages from 8 weeks prior to TMI to over 14 weeks after TMI. Simultaneously the mean time between refuelings increased from approximately 12 months prior to TMI to nearly 18 months after TMI. Using monthly data on U.S. NPPs in the post-TMI era we estimate parameters of the utility's profit function, the failure processes that lead to unplanned forced outages, and the parameters governing the duration of refueling outages. These parameters imply an endogenous distribution of operating spells and capacity utilization levels that depend on the NPP's age, signals the operator receives about the NPP's current operating state, and the duration since last refueling. The estimates of the DP model reveal that utilities appear responsive to NRC regulation insofar as they impute a very high cost to unplanned and forced outages. Utilities are also highly averse to causing unnecessary wear and tear on their NPP's caused by stop/start operation of their NPP's including planned and unplanned outages. Overall, the DP model yields very accurate predictions of nuclear power generation including the impact of the relatively rare event of NPP decommissionings.[1]

JEL Classification: C41–Duration Analysis

Keywords: dynamic programming, electricity generation, nuclear power plants

---

[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 grants provided indirect support for this research project.

1. Introduction

Much attention has been paid to the collapse of the nuclear power industry in the U.S. following the Three Mile Island accident in March 1979 (see, e.g., Campbell, 1988). Due to substantial increases in construction costs, leadtimes, 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 over 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). Despite these problems, the nuclear power industry is significant producer of electricity: in 1992 NPPs with a combined generating capacity of 100 gigawatts (about 15% of the country's total capacity) produced 620 billion kilowatt-hours, with an average capacity factor of 71%, or 22% of the nation's electricity supply (INPO, 1993).

Increasing recognition of the adverse environmental impacts of fossil fuels, including acid rain and the "greenhouse effect," has lead to a reconsideration of the nuclear energy option (see, e.g., Cohen, 1990, and Rothwell, 1994). Although demand for electricity is projected to grow at annual rates of approximately 2% over the remainder of the 1990's (with peak power demand to be about 650 gigawatts in 1999), there will be a significant loss of generating capacity over the next two decades due to retirements of aging NPPs. The U.S. Nuclear Regulatory Commission (NRC) issues 40-year operating licenses to NPPs. The NRC estimates that if licenses of existing plants are not extended, the loss in generating capacity due to retirements will be about 6 gigawatts per year from 2005 to 2010. According to some estimates, extending the lives of existing NPPs by 20 years could save U.S. consumers about $350 billion (Forrest and Deutsch, 1988). Although the NRC is currently preparing procedures for extending operating licenses for plants that can pass its safety standards, a variety of age-related problems, such as radiation-induced embrittlement of the reactor vessel, have lead to the presumption that license extension would not be economic. It is not clear whether utilities would find it profitable to continue operating existing NPPs for 40 years. During the last seven years, five NPPs have been closed more than ten years before their license expiration dates: La Crosse (after 19 years of operation), Fort St. Vrain (after 20 years), Rancho Seco (after 24 years), Yankee Rowe (after 30 years), and San Onofre 1 (after 25 years).

Forecasting nuclear energy's future contribution to the U.S. electricity supply requires a model of NPP operations that can predict output levels and probable dates of retirement. This paper presents and estimates a dynamic programming model of utility's optimal operating policy for NPPs that includes the option of closing unprofitable NPPs. Our model provides accurate predictions of nuclear power generation including the rare event of an NPP closure.

The electricity generated by an NPP equals the product of its net electrical generating capacity (as measured by its net "maximum dependable capacity" or MDC), and its *capacity factor,* the ratio of the electricity the plant generated over a period to the amount it would have generated if it had been running continuously at its MDC. NPP capacity factors in the U.S. have increased from 59.8% in 1980 to 66.1% in 1990. [2] The capacity factor can be further decomposed into the product of the plant's capacity utilization rate and its *availability factor.*[3] The availability factor is the percentage of the time that the NPP is generating electricity and is determined by two variables: the plant's planned availability (equal to 1 minus the fraction of time spent in planned outages for refueling and preventive maintenance) and the plant's reliability, as measured by the fraction of time spent in unplanned or forced outages.

During planned refueling outages the NPP is shut down, the reactor head and control rod assembly are removed, and older fuel rods are replaced with new ones. Most of the major preventive maintenance activities, including inspection and maintenance of "hot areas" such as the reactor core, cooling system, and steam generators, occur during this time. Plants occasionally experience even longer outages associated with major capital upgrades, such as NRC-mandated retrofits, or replacement of age-degraded reactor components, such as nuclear steam generator in a pressurized water reactor. The NRC can impose a shutdown until safety concerns are resolved. Most of the "major problem" outages last many months. Some of them, such as the shutdown of the Browns Ferry NPPs, can last several years.

There is a trade-off between safety and cost-efficiency in NPP operations: longer or more frequent maintenance and refueling outages improve reactor safety, but are expensive. For a typical 1,000 megawatt NPP, the opportunity cost is roughly $1 million for each day of lost output. Due of the huge social costs of a catastrophic nuclear accident and the potential conflict of interest between profitable operation and public safety, the nuclear industry is subject to extremely exacting and costly regulation by the NRC, which has the authority to shut down plants, levy fines, mandate equipment retrofits, and conduct surprise inspections.

NRC regulation has been especially intense in the aftermath of the Three Mile Island (TMI) accident. The accident, due to an unfortunate combination of equipment malfunctions and operator errors, resulted in a partial core meltdown, a release of radioactive steam from the containment building, and over $4 billion in cleanup costs, liability and litigation. [4] The increased stringency is reflected in the frequency and magnitude of NRC fines. These increased

[2] See EIA (1992, p. 54). These are low in comparison to Canadian heavy water reactors, CANDUs. These NPPs have higher capacity factors (nearly 80%) partly because they do not need to be shut down for refueling, unlike U.S. light water reactors.

[3] The availability factor is also known in the industry as the NPP's *service factor.* An NPP that is operating may not always be running at full capacity, i.e., generating electricity at its MDC. Reasons why operating NPPs may choose to operate at less than MDC include temporary power reductions due to NRC-mandated deratings and output reductions designed to "follow demand" when the plant is feeding a power grid that is near its capacity. For a more in depth analysis of these issues see Rothwell (1990).

[4] Before TMI there was a general impression that the NRC was lax in its imposition of regulations. However after the Kemeny Commission report there was a shake-up in the NRC, leading to much more rigorous regulation of the industry: "I believe the Three Mile Island accident formally and finally broke the ties connecting the NRC to the development of nuclear power, which was part of its heritage from the AEC. I will not argue the issue as to how much of a link existed in the regulatory side of the AEC or whether it was a major influence in the early years of the NRC. Certainly many accused the NRC of such a link. The Three Mile Island accident has severed any such ties. The Three Mile Island accident was a catharsis. In Greek tragedy, such a catharsis must occur, for a purging. The Three Mile Island accident was the NRC's catharsis, and it now becomes a regulatory agency." John Ahearns, NRC Commissioner, 1981.

from a rate of 0.07 per plant with an average penalty of \$25,000 in 1979 to a rate of 1 per plant and an average penalty of \$84,000 in 1989 (penalties in 1984 dollars). Following the accident, the NRC introduced its TMI Action Plan. It mandated changes in operator training, retrofits of control room instrumentation and certain components of the NPP cooling system, increased plant inspections and reporting requirements, and unprecedented multi-year shutdowns of plants. See Dubin and Rothwell (1989).

The NRC's stricter regulations appear to have had a payoff in terms of increased safety: according to the National Research Council (1992) the rate of unplanned automatic scrams[5] decreased from an average rate of 7.4 per plant in 1980 to 1.6 in 1990, and the rate of unplanned safety system actuations decreased from 1.3 per plant in 1985 to 0.7 in 1990. David, Maude-Griffin, and Rothwell (1994) used a Cox proportional hazard model to study the durations between successive unplanned outages of a sample of U.S. NPPs. They found that there were substantial reductions in the risk of unplanned outages after TMI: the mean duration between unplanned outages more than doubled from 26 days before TMI to 53 days after TMI. They concluded, "The estimates suggest that there were substantial reductions in the risk of unplanned outage after TMI, that these reductions were concentrated in the plants which had the highest levels of unplanned outage risk prior to TMI, and that most of the risk reduction occurred several years after the TMI accident." (p. 1).

However the NRC's stricter safety regulations have come a substantial cost in reduced productivity and increased power generation costs. Rahn, *et al* (1984, Section 19.3) estimated that in 1979 NRC-mandated shutdowns resulted in a loss of 12% of industry electrical generation capacity, or about 30 billion kilowatt-hours of potential output. Besides shutdowns, NRC regulations contributed to a decline in availability factors by more than doubling the median duration of refueling outages from approximately 8 weeks in the period 1974-1978 to over 18 weeks in 1982.[6] The industry has claimed that these increased durations were largely the result of additional NRC-imposed surveillance, backfit, and maintenance requirements (National Research Council, 1992, p. 51). Utilities have responded to these increased refueling durations by increasing the mean duration between refuelings from 12 to 18 months. These increases have partially counteracted the decline in capacity factors caused by longer refueling durations.

The reason longer cycles improve capacity factors is straightforward: abstracting from unplanned shutdowns and load following, a plant with an operating duration of O months and a refuel duration of R months will have a long run capacity factor of $O/(O+R)$, which is monotonically increasing in operating duration, O. For example, if refuelings occur annually and last 2 months, the plant's maximum capacity factor (assuming no forced outages) is

[5] The term "scram" refers to the rapid, typically automatic insertion of the reactor control rods into the reactor core to halt nuclear fission and prevent the core from overheating.

[6] See Stoller (1989, Section 4). The greatest force of the stricter NRC regulations seems to have occurred in the 4 or 5 years following the TMI accident: in the period 1986-88 refueling durations decreased to a median of about 13 weeks. We confirm this in our analysis in section 3.

84%. By moving to 24 months between refuelings the maximum capacity factor improves to 92%. These calculations assume that the duration of a refueling outage does not increase with the length of the previous operating cycle. In Section 5 we show that this is a good assumption. An implication of this simple formula for the capacity factor is that the potential gain to moving to longer operating cycles is greater the longer the mean refueling duration R. For example, if $R = 3$ months, the increase in capacity factor of moving to 18 month operating cycles is 7.7% versus only 3.8% if $R = 1.5$ months. This factor may explain a finding in the study by Stoller (1987): "The very best performers have remained on short cycle for a variety of reasons, one of which being that they have a smaller incentive to change to long cycles." (p. 2-12).

There are two important tradeoffs in moving to longer operating cycles: (1) longer cycles involve a higher cost of nuclear fuel per unit of time, and (2) longer cycles could involve higher rates of planned and unplanned outages toward the middle and end of an operating spell. Fuel costs increase during long cycles for at least two reasons: first, long cycles require higher quantities or higher enrichment levels in the nuclear fuel; and second, fuel efficiency is lower the longer the operating cycle. In sections 3 and 5 we show that the forced outage rate increases with the duration since last refueling, at least after 12 months of operation. The utility needs to balance the benefits of less frequent refueling outages against the increase in fuel costs and the increased rate of planned and unplanned outages associated with longer operating cycles, especially toward the end of an operating spell.

The increased stringency of NRC regulations after TMI are reflected in the rapid escalation in real operating and maintenance (O&M) costs of U.S. NPPs: these costs increased at a rate of 14% per year between 1974 and 1984 (see EIA, 1991, p. 22). The equivalent costs for fossil fueled plants decreased from 1982 to 1987 (National Research Council, 1992). EIA (1991) concluded that tighter NRC regulation in the post-TMI era was a major factor behind the operating cost escalation, responsible for over 50% of the increase in operating and maintenance costs and most of the increase in the cost of capital additions (e.g., backfits) from 1974 to 1987. While its tempting to place the blame for the financial woes of the nuclear industry on the NRC's stricter regulations following TMI, it is important to keep in mind that there is a wide variation in NPP reliability and operating costs. David, Maude-Griffin, and Rothwell (1994) found "One of the more remarkable features of our results is the degree of variation in the estimated unplanned outage risk across plants." (p. 17). But they found it was difficult to characterize factors that correlate well with differences in outage risk: "It is easy to find instances where plants of similar size and age with the same NPP vendor are at opposite ends of the spectrum of risk estimates." (p. 17). There is similar variation in O&M costs: EIA (1991, Table 2) shows annual nonfuel O&M costs ranging from \$30.26 to \$158.08 per kilowatt of capacity. The best performing NPPs can be profitable, generating electricity at well below the cost of the most efficient fossil-fueled plants. For example, EIA (1992, Table 6) shows that in 1990 the most efficient NPPs (i.e., the plants in the lowest quartile in the distribution of total generation expenses) produced electricity at 1.54 cents per kilowatt-hour compared to 1.89 cents per kilowatt-hour for the most efficient quartile of coal plants. However, the cost disadvantage of nuclear power becomes clear when we compare the median generation costs: 4.69 cents per kilowatt-hour for NPPs versus 2.76 cents per kilowatt-hour for coal.[7]

Besides considerable variation in performance across different plants in the same year, our analysis in Section 3 shows that there is considerable variation in the performance of a given plant at different points in time. This variation reflects the considerable intertemporal risk of unplanned outages. This implies that one can get a very different picture of the operating efficiency of an NPP depending on the sample frame over which one observes its operation. The intertemporal variability in performance complicates the utility's decisions about whether a specific plant is or is not profitable. However even if an NPP has been clearly identified as inefficient in the sense that its cost of power generation is systematically higher than a fossil-fueled plant of equivalent capacity, it still does not follow that it is necessarily optimal to close the plant for decommissioning. There are important reasons why a utility might want to continue operating an "inefficient" NPP: (1) an NPP involves substantial unrecoverable sunk costs since most of the equipment cannot be converted for use in other generating facilities; and (2) closing and decommissioning an NPP is an extremely expensive procedure. Decommissioning a large scale NPP could cost hundreds of millions of dollars. Given that the utility does not have the option of costlessly transforming its NPP into a fossil fueled plant, or even of costlessly disposing of its NPP, it follows that it may indeed be optimal to continue operation. See Pasqualetti and Rothwell (1991).

Some of the NPPs that have been closed were smaller reactors, such as the La Crosse plant, or had "experimental" designs, such as the high-temperature, gas-cooled reactor at Fort St. Vrain. Other closures have occurred after the discovery of "major problems" that would have involved NRC-mandated shutdowns until the problems were corrected. An example is the discovery of small cracks in the reactor vessel at Yankee Rowe in 1991. The utility decided to close Yankee Rowe in 1992 because its only feasible option to resume operation was to replace the reactor vessel.[8] The utility completed decommissioning in 1994 before the low-level waste repository at Barnwell, South Carolina, closed its doors to out-of-region waste. In some cases plant closures may not be strictly voluntary decisions on the part of the utility. For example, the Rancho Seco plant was closed after the local utility lost a public referendum that forced it to close. In contrast, our dynamic programming model assumes that all plant closures are voluntary decisions by the utility. However the occurrence of long "major problem" shutdowns are treated as random events that are not voluntarily initiated by the operator. Some of these can be due to political or regulatory problems, as with

[7] For the least efficient quartile, total generation costs were 40.21 cents per kilowatt-hour for nuclear power versus 28.55 cents per kilowatt-hour for coal.

[8] If an embrittled reactor vessel is detected before the formation of cracks, it is possible to partially "regenerate" the reactor vessel by an expensive thermal annealing operation that restores the vessel's ductility. However thermal annealing cannot repair cracks in the reactor vessel, in which case the utility is faced with the alternative of replacing the reactor vessel or closing the plant.

Rancho Seco. After the plant has involuntarily entered a major problem spell it is much more likely that it will take the additional voluntary action of closing the plant for decommissioning.[9]

The remainder of this paper explores these issues and utility behavior in more depth. Section 2 provides a brief overview of some basic concepts of nuclear power generation and a description of operating and fuel management strategies employed at U.S. NPPs. Section 3 describes the data used for this study, derived from the NRC's "Graybooks" from 1975 to 1993. Using this dataset, we confirm the above noted changes in durations of refueling and operating spells and in the risk of forced outages that have occurred after the TMI accident.

Section 4 presents our dynamic programming (DP) model of optimal operation of an NPP. It builds on previous work by Sturm (1993), which focused on operations of European NPPs. In the DP model the NPP operator must decide each month whether to operate, refuel, or close the plant. If the operator decides to run the plant, there is an additional (continuous) decision of how much power to produce, which we model as the operator's capacity utilization decision chosen from the $[0, 1]$ interval. The operating decisions depend on the signals the operator receives about the NPP's operating status, some of which are recorded in the NRC's Graybook data and some of which are unobserved by the econometrician. Our DP model is designed to accommodate both types of signals. To simplify computation of the DP model we approximate the operator's continuous utilization decision by discretizing the $[0, 1]$ interval into six utilization subintervals.[10] In Section 5 we use monthly data on 111 U.S. NPPs in the post-TMI era to estimate the unknown parameters of the utility's profit function, the failure processes that lead to unplanned forced outages, and the parameters governing the duration of refueling outages.[11] These parameters imply an endogenous distribution of operating spells and capacity utilization levels that depend on the NPP's age, signals the operator receives about the NPP's current operating state, and the duration since the last refueling. The estimates of the DP model reveal that utilities impute a high cost to unplanned and forced outages. Utilities are also highly averse to causing unnecessary wear and tear caused by stop/start operation of their NPPs, including planned and unplanned scrams. Section 5 presents an analysis of the goodness of fit of the DP model, comparing predictions of the DP model to the actual operating behavior of U.S. NPPs. Generally, the DP model yields accurate predictions of nuclear power generation including the impact of the rare event of an NPP closure. Section 6 presents concluding remarks and directions for future research.

---

[9] This way of viewing plant closures might not be a bad approximation since the utility that operated Rancho Seco had the option of "mothballing" the plant until a more favorable political climate emerged that would allow it to resume operation. In this sense, the closure decision was just as "voluntary" as the decision to close Yankee Rowe: in both cases substantial costs would have to be incurred that were deemed too high in relation to the short remaining licensed lives of these NPPs.

[10] The approximation error involved in making this discretization is negligible since during most months utilization rates are either 0 or 1: intermediate load factors are typically chosen so that the NPP can be temporarily shut down to investigate problems.

[11] Our sample from 1975 through 1993 has 116 NPPs. Two of these were closed before 1980 and five were closed after 1980. Therefore, our sample size varies from year-to-year.

2. Nuclear Power Technology and Operating Practices at U.S. Nuclear Power Plants

This section provides a brief overview of nuclear power generation and discusses the trade-offs involved in operating an NPP, including the important topic of fuel management and scheduling refueling outages. Some background on these issues is necessary in order to understand our specification of the DP model in Section 4 and to interpret the estimation results in Section 5.

There are many types of NPPs in the world, but in the U.S. nearly all commercial NPPs use either pressurized water reactors (PWRs) or boiling water reactors (BWRs). Of the 111 licensed U.S. NPPs operating in 1993, 76 were PWRs and the remaining 35 were BWRs.[12] Westinghouse manufactured 44 of the PWRs, Combustion Engineering produced 15, and Babcock & Wilcox made nine (ten if we include the ill-fated Three Mile Island 2 reactor). All BWRs in the U.S. have been designed and built by General Electric. PWRs and BWRs are types of light water reactors (LWRs) using ordinary water as coolant and moderator.[13] Figure 2.1 (taken from Nero 1979) illustrates the placement of the reactor within the steel-lined, reinforced concrete *containment building* designed to prevent release of radioactivity into the environment in case of an accident. The figure shows a crane within the containment building which is used to remove a protective shield covering the reactor and the reactor head, control rods, and fuel elements during a refueling. The containment building in figure 2.1 houses a PWR since it also contains *steam generators* that convert superheated water from the reactor into steam used to drive the turbines (which in turn drive the electrical generators). A typical PWR consists of two cooling loops as illustrated in Figure 2.2 (also taken from Nero, 1979). In the inner cooling loop superheated water at temperatures between 550 to 600 degrees F is pumped through the reactor core at a pressure of about 2200 pounds per square inch at a rate of 140 million pounds per hour.[14] The flow of water removes heat produced by the reactor (whose interior temperatures are over 4,000 degrees F) and transfers this heat to the secondary cooling loop via one or more *steam generators* that produce the steam that drives turbine generators that produce electricity. The overall thermal efficiency of a PWR (i.e., the ratio of the net electrical energy produced to the total thermal energy generated) is about 32%. A BWR, on the other hand, has only a single cooling loop. As the coolant passes over the reactor core it is allowed to boil, directly producing the steam the runs the turbine generators. The BWR design allows the coolant to circulate under lower pressure (about 1,000 pounds per square inch), and the elimination of the steam generators allows BWRs to run at slightly higher thermal efficiencies, approximately 34%. The disadvantage, however, is that the boiling of the coolant in the reactor core results in lower coolant densities

---

[12] The exception is the high-temperature, gas-cooled NPP at Fort St. Vrain, Colorado, which was closed in 1989.

[13] LWRs use ordinary (light) water to both cool the reactor and control, i.e., moderate, the nuclear reaction. The Canadian CANDU uses heavy water that absorbs fewer neutrons than light water, allowing it to use unenriched uranium as fuel.

[14] It takes a great deal of energy to move this much coolant through the core at these pressures: the energy required to drive the pumps for the inner cooling loop consumes about 5% of the reactor's gross electrical output.

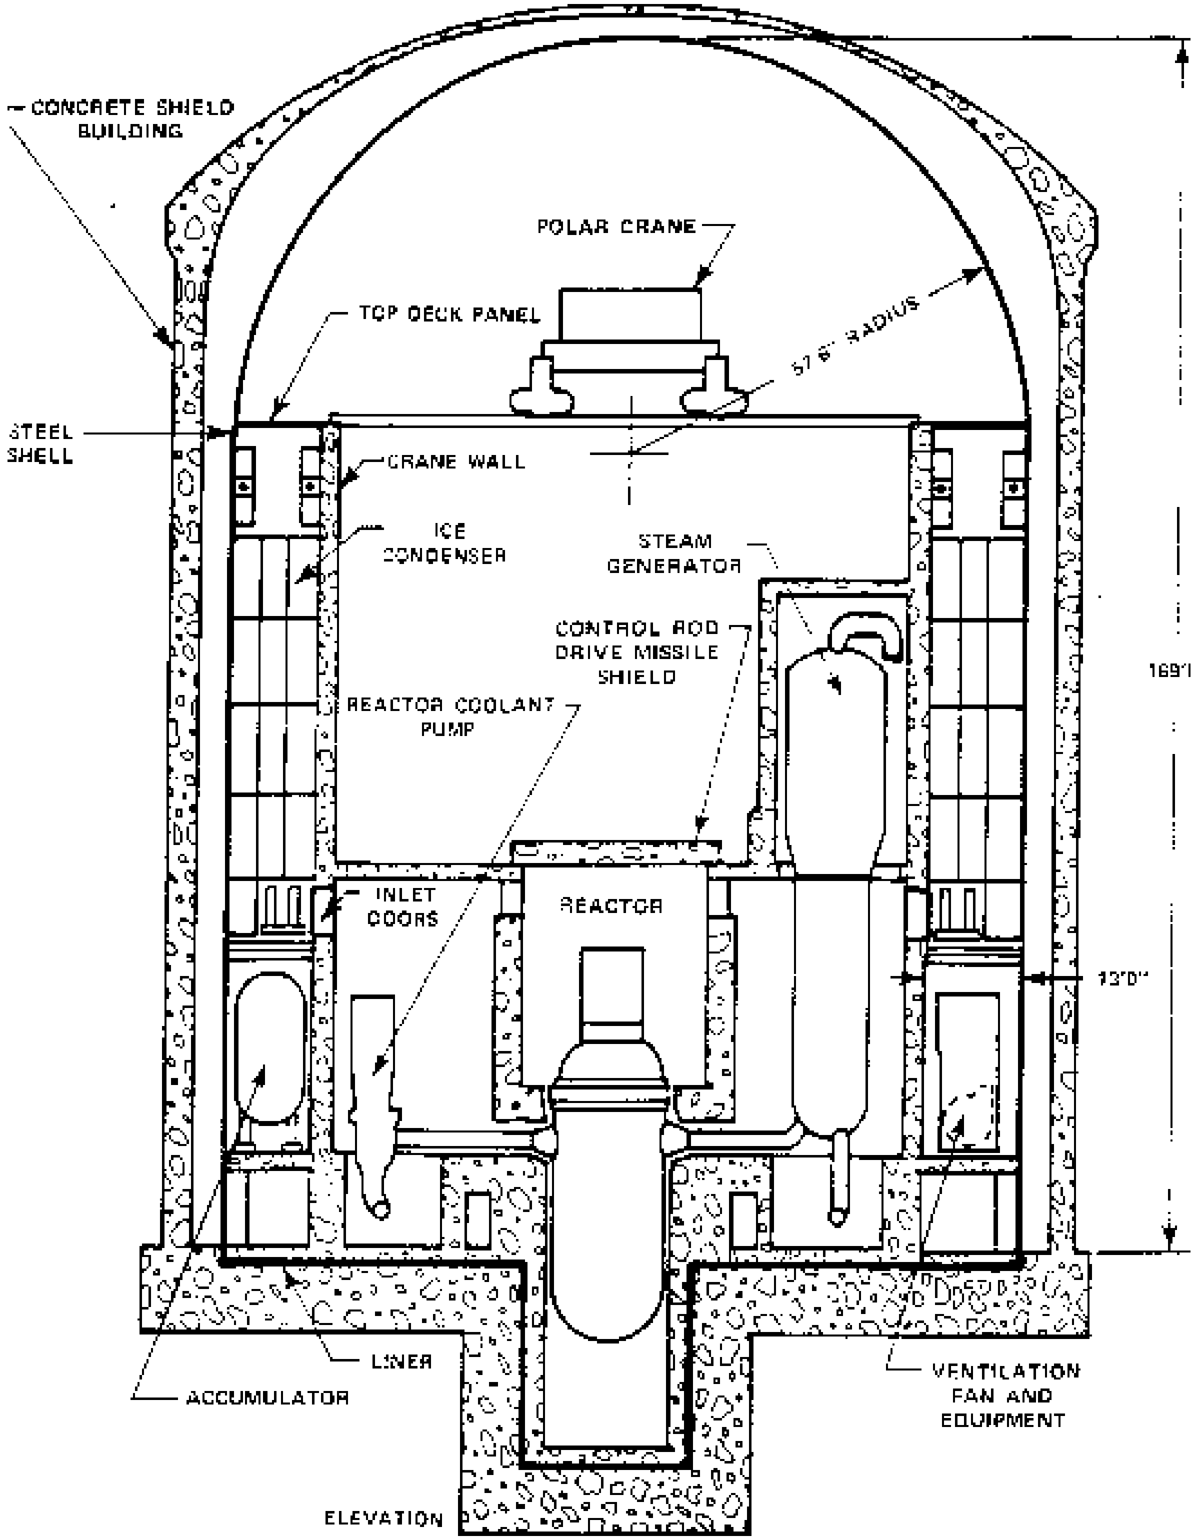

Figure 2.1 Containment Building for a Pressurized Water Reactor

making it more difficult to moderate power generation in the core than in PWRs. In addition, the steam driving the turbine generators is slightly radioactive since it comes from water passing through the reactor core.

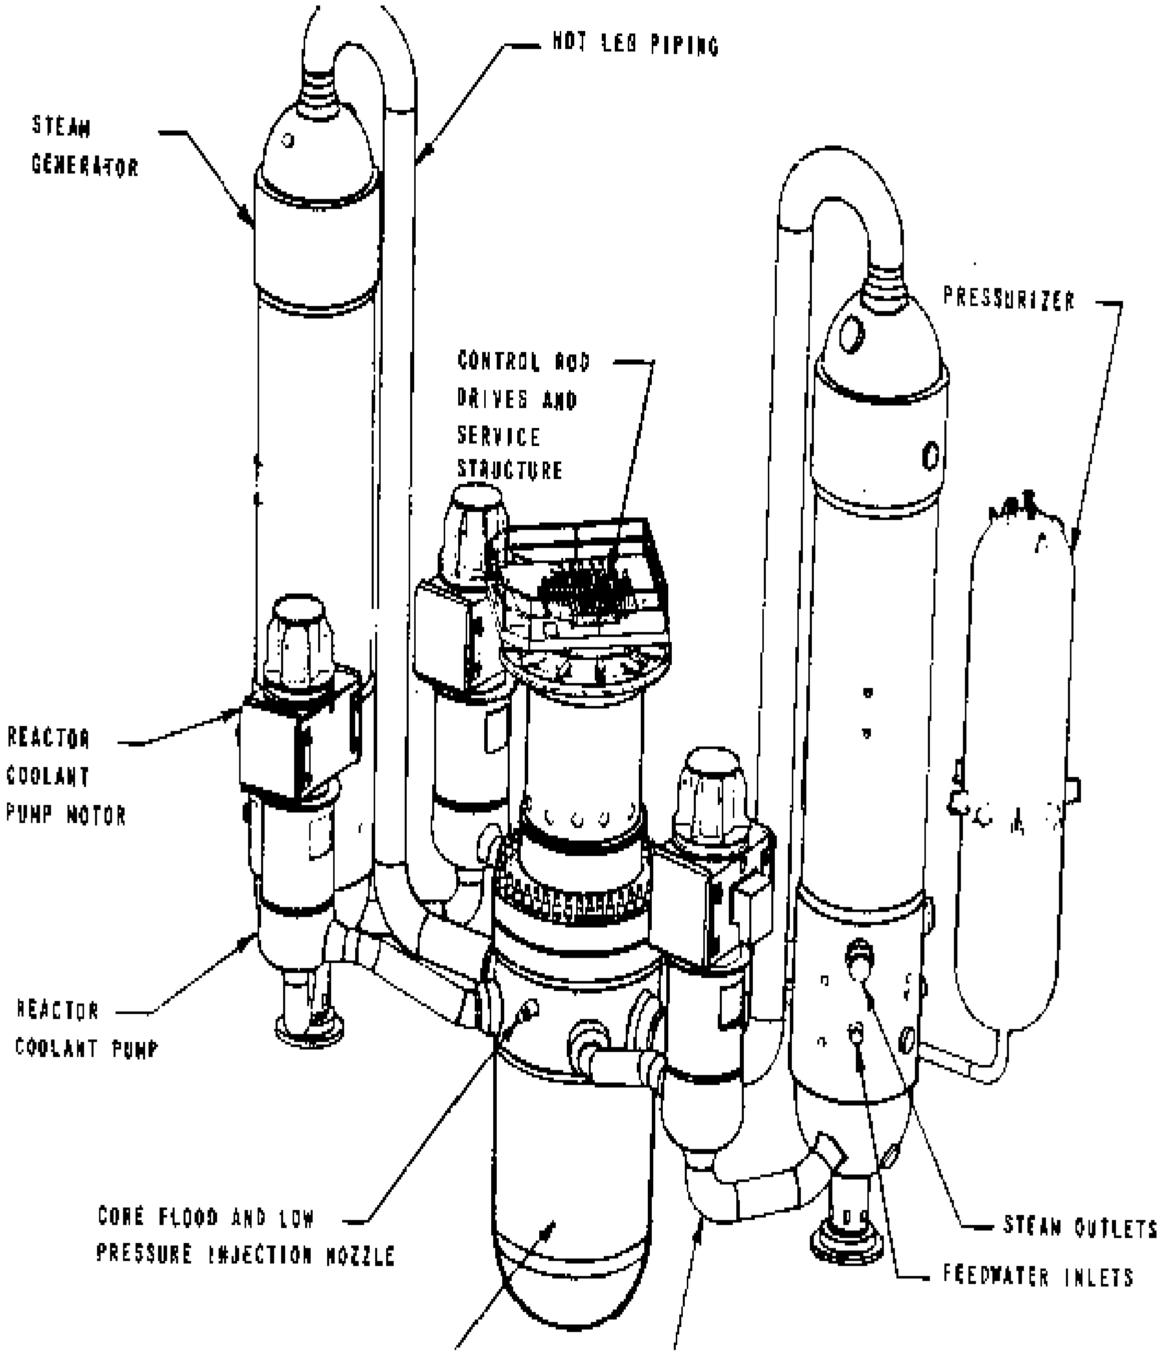

Figure 2.2 Main Components of a Pressurized Water Reactor

NPPs generate power via nuclear fission using slightly enriched uranium as fuel. Naturally occurring uranium is composed of two isotopes: fissile U^{235} (0.7% by weight) and nonfissile U^{238} (99.3% by weight). During a nuclear reaction, U^{235} releases neutrons that can split other U^{235} atoms or be absorbed by U^{238} atoms. Fission releases tremendous energies: a single fission of an U^{235} atom releases 192 million electron volts of energy, so just one kilogram of U^{235} has the same energy content as 3,000 tons of coal. A sufficient concentration of U^{235} in the core of a reactor can start a *chain reaction*. The minimal fissile material needed to sustain a chain reaction is known as a *critical mass*. The level of reactivity is denoted by the symbol k: it is the ratio of neutrons born in generation t to the neutrons born in generation $t-1$. When $k=1$, the reactor is said to be *critical* and the chain reaction is in equilibrium and can be sustained indefinitely. If $k>1$, the reactor is *super-critical* and the neutron population and power levels increase exponentially until k is reduced to 1 with the introduction of neutron-absorbing moderators, such as boron in the control rods or in the coolant flow. If $k<1$, the reactor is *sub-critical* and the chain reaction will die away unless an independent source of neutrons is introduced. The basic principle of operating an NPP is to maintain the reactor at the critical level via continuous adjustment of control rods and adjustment of concentration of moderators in the reactor coolant stream. Natural uranium does not have a sufficient concentration of U^{235} to form a critical mass in a LWR, so it must be enriched to a concentration of 2-4% U^{235} by weight. Starting with an initial enrichment of 3.2% U^{235}, fuel discharged from the NPP after usual exposure will have decomposed into 0.85% U^{235}. About 3.8% of spent fuel consists of fission products, some of which are strong moderators.

A 1,000 MWe NPP consumes about 30 tons of nuclear fuel per year, generating an equal amount of high-level radioactive waste. The high kinetic energy released by fission has deleterious effects on the structure of the fuel. Some fission products, such as xenon, appear as gases that eventually create substantial pressure within the fuel. As a result, the fuel 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 fissile U^{235} and buildup of fission products, combined with the effect of radiation-induced fuel swelling and distortion, are the 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 fuel, the design of the fuel rods, and the fuel management strategy adopted by the operator.

During refueling, the reactor head and control rod drive shaft are removed, exposing the tops of the fuel elements inside the reactor core.[15] Due to the exposure of the core, plant workers can receive a substantial radiation dose during a refueling. A standard practice in controlling this dose is to flood a large tank containing the reactor vessel and exchange the fuel elements underwater. The spent fuel elements are moved into an auxiliary fuel storage pond from an opening inside the containment building and are replaced with fresh fuel assemblies. Most NPPs adopt *batch refueling* strategies where typically only one third to one fourth of the fuel assemblies are replaced during a

[15] In a BWR control rods are inserted from the bottom of the reactor, although fuel assemblies are removed from the top of the reactor where the steam outlets are located.

refueling. Partially used fuel assemblies are rotated to different parts of the reactor core, a topic we will return to in our discussion of fueling strategies below.

Reactor vessels are massive: a typical vessel is about 40 feet high and 15 feet in diameter and 7 to 8 inches thick and can weigh as much as 450 tons. This size is necessary to withstand the high temperature, high pressure, and constant neutron bombardment. Figure 2.3 (taken from Nero 1979) presents a cut away diagram of a BWR, showing the placement of the fuel elements, control rods and radiation shielding within the reactor vessel. The reactor core of a typical PWR contains approximately 200 *fuel assemblies* containing arrays of enriched uranium fuel rods, illustrated in figure 2.4.[16] A typical PWR contains approximately 40,000 fuel rods. The fuel rods contain a series of cylindrical uranium pellets (about one-half inch long and one-half inch in diameter) containing uranium dioxide enriched to a concentration of 3-5% U^{235}, encased in a corrosion-resistant cladding, usually made of zirconium. The core of a 1,000 megawatt PWR contains about 100 tons of uranium dioxide, and each fuel assembly typically releases the energy equivalent of 250,000 barrels of oil during its occupancy in the reactor core. Interspersed in the fuel assembly's array of fuel rods are control rods of a neutron-absorbing alloy encased in stainless steel tubes. Also, one or more tubes in the fuel assembly are reserved for monitoring instruments. In a PWR approximately one-third of the fuel assemblies contain one to two dozen control rods. A typical BWR has many more fuel assemblies interspersed with regularly spaced control rods. The control rods are the primary means of reactor moderation. Rapid insertion of the rods into the reactor core (known as a scram) will bring the chain reaction to a halt, reducing the power generation of the reactor to 6.5% of its steady-state level within one second after insertion.[17] All commercial NPPs in the U.S. have extensive monitoring systems that will trigger an automatic scram upon detection of abnormal conditions.

The operator's basic objective is to keep the reactor critical (at $k = 1$) and ensure safe operation in all circumstances, including unexpected events, such as equipment failures, or reductions in turbine load due to variations in electrical demand. A secondary objective is to optimize fuel utilization. The level of fuel utilization in an NPP is its *burnup* rate. It is measured in megawatt days of energy released per ton (MWd/t) of uranium fuel in the reactor core. Increasing the burnup rate reduces the uranium fuel input and the radioactive waste output. A final objective in NPP operation is to obtain an even power distribution in the reactor core. If the power distribution is uneven, fuel must be replaced when the most exposed part has reached its burnup limit while the rest of the fuel has yet to deliver its maximum potential.

From an economic point of view, the ideal pattern of operation would be to run at full capacity for as long as possible and have refuelings that are as short as possible. However as noted above, fuel integrity is the key factor

---

[16] This figure was taken from Pershagen, 1989.

[17] There is considerable residual heat from continued radioactive decay in the fuel elements even a year after closure. As a result, it is necessary to maintain continuous coolant flow within the reactor core after shutdown.

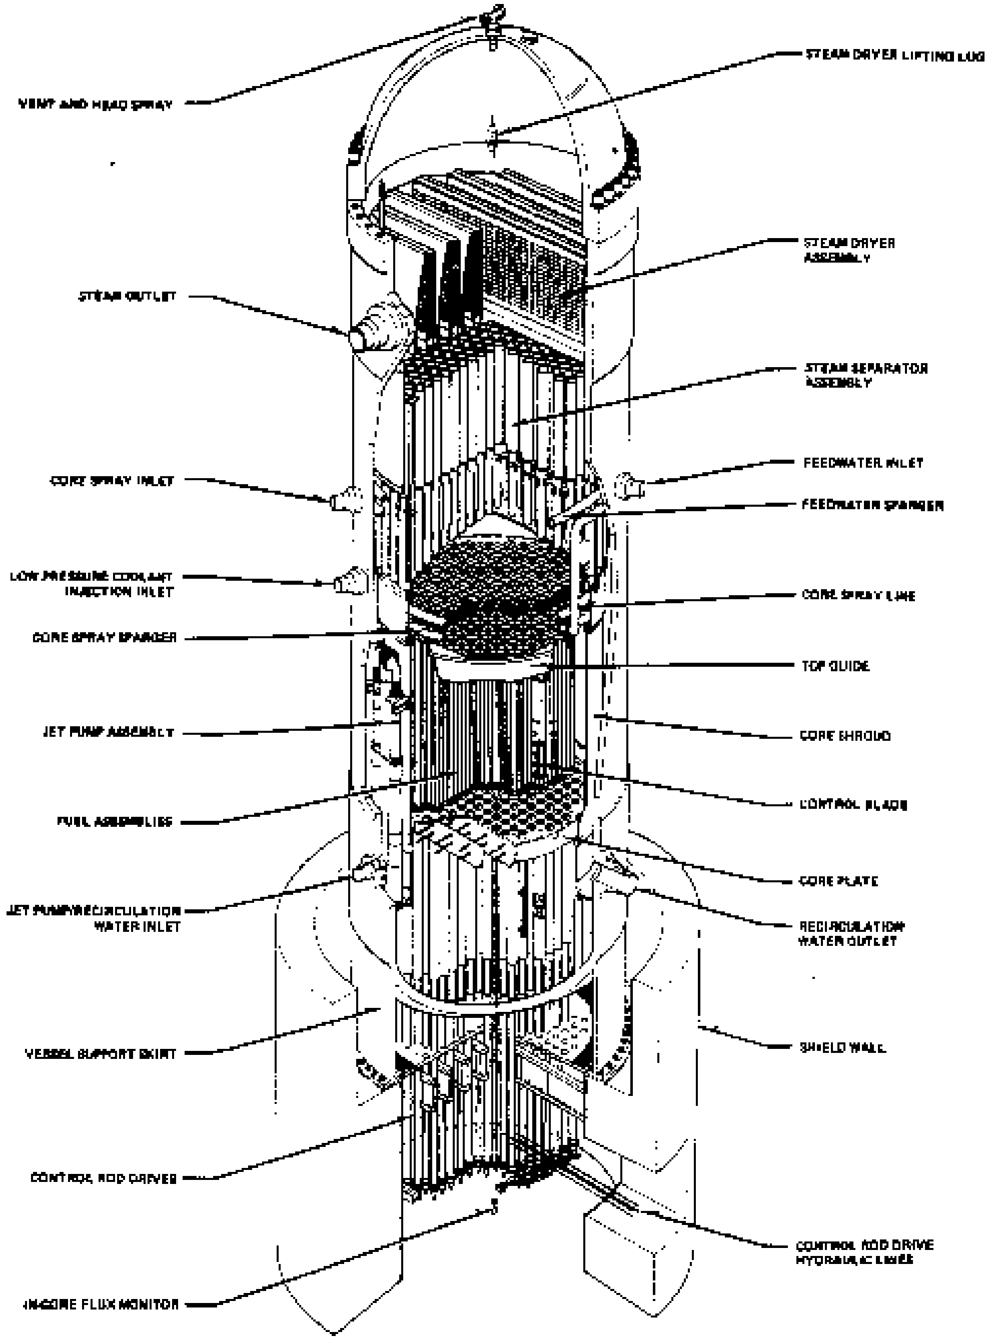

Figure 2.3 Cut Away Diagram of Internals of a Boiling Water Reactor

limiting the maximum duration between refuelings. Distortions in uranium fuel pellets caused by release of gases can lead to swelling or cracking of the fuel cladding and release of radioactive isotopes into the reactor coolant. This should be avoided for two reasons: (1) radioactivity would be deposited throughout the system making maintenance more difficult and (2) fuel cladding integrity is a major barrier to the release of radioactivity should there be an accident. If more than a few of the 40,000 fuel rods fail during an operating spell, the NPP must be shut down and the failing fuel assemblies removed. Although failure rates of fewer than 5 rods per operating cycle have been achieved in LWRs, the NRC closely monitors fuel failure and excessively high failure rates can lead to a derating (i.e., a reduction in the maximum allowed power generation). In BWRs fuel failures are the largest single cause of plant unavailability next to refueling, accounting for approximately 750 hours per year in reactor downtime. Since the length of an operating cycle is limited by the maximum permissible burnup of the oldest fuel rods within the reactor core, technological improvements in “extended burnup” fuel rods have been a key factor in enabling NPPs to move to longer operating cycles. Before the 1980s, fuels were designed for an average burnup of 35,000 MWd/t, but improvements in fuel technology have reduced failure rates and enabled maximum burnups as high as 60,000 MWd/t.

In addition to the tradeoff between the length of the operating cycle and fuel reliability, there is also a tradeoff between the length of the operating cycle and fuel efficiency. Longer operating cycles generally entail higher fuel expenses per unit of time as illustrated in Figure 2.5 (taken from Rahn, *et al* 1984). Assuming burnup remains constant, moving to a longer operating cycle requires more fuel (i.e., more the fuel assemblies are replaced at each fueling) or a higher level of enrichment, or both, since at each refueling the operator must load the core with sufficient reactivity to maintain criticality until the next refueling. The excess neutron production at the beginning of the operating cycle must be absorbed by increased levels of moderators, limiting the potential thermal generation. By minimizing the number of neutrons absorbed by reactor coolant, moderators, and control rods, the operator maximizes the neutrons available for energy producing purposes.

The higher the initial reactivity of the fuel, the more problems the NPP operator has with uneven burnup and underuse of the fuel. To ensure more even burnup, special fuel rotation strategies are employed, including *out-in refueling* (in which fresh fuel is placed in the periphery and used fuel is rotated to the center) and *scatter refueling* (in which new fuel assemblies are interspersed with used fuel assemblies). If the nuclear fuel and moderators were uniformly distributed in the core, the neutron density would be highest in the center of the reactor core and lowest in the periphery where more neutrons would be absorbed by the reactor’s heat shield. Refueling would be required when the fuel rods in the center of the core reached their maximum permissible burnups, but fuel rods at the periphery would not have released their full energy potential. Thus the engineering objective of maximizing fuel burnup argues for more frequent *partial batch refuelings* in which only a small fraction of the fuel assemblies in the center of the core are removed, partially burned assemblies at the periphery are moved to the center positions in the core, and new fuel assemblies are added at the periphery. In the limit, the ideal strategy for maximizing burnup would be to refuel

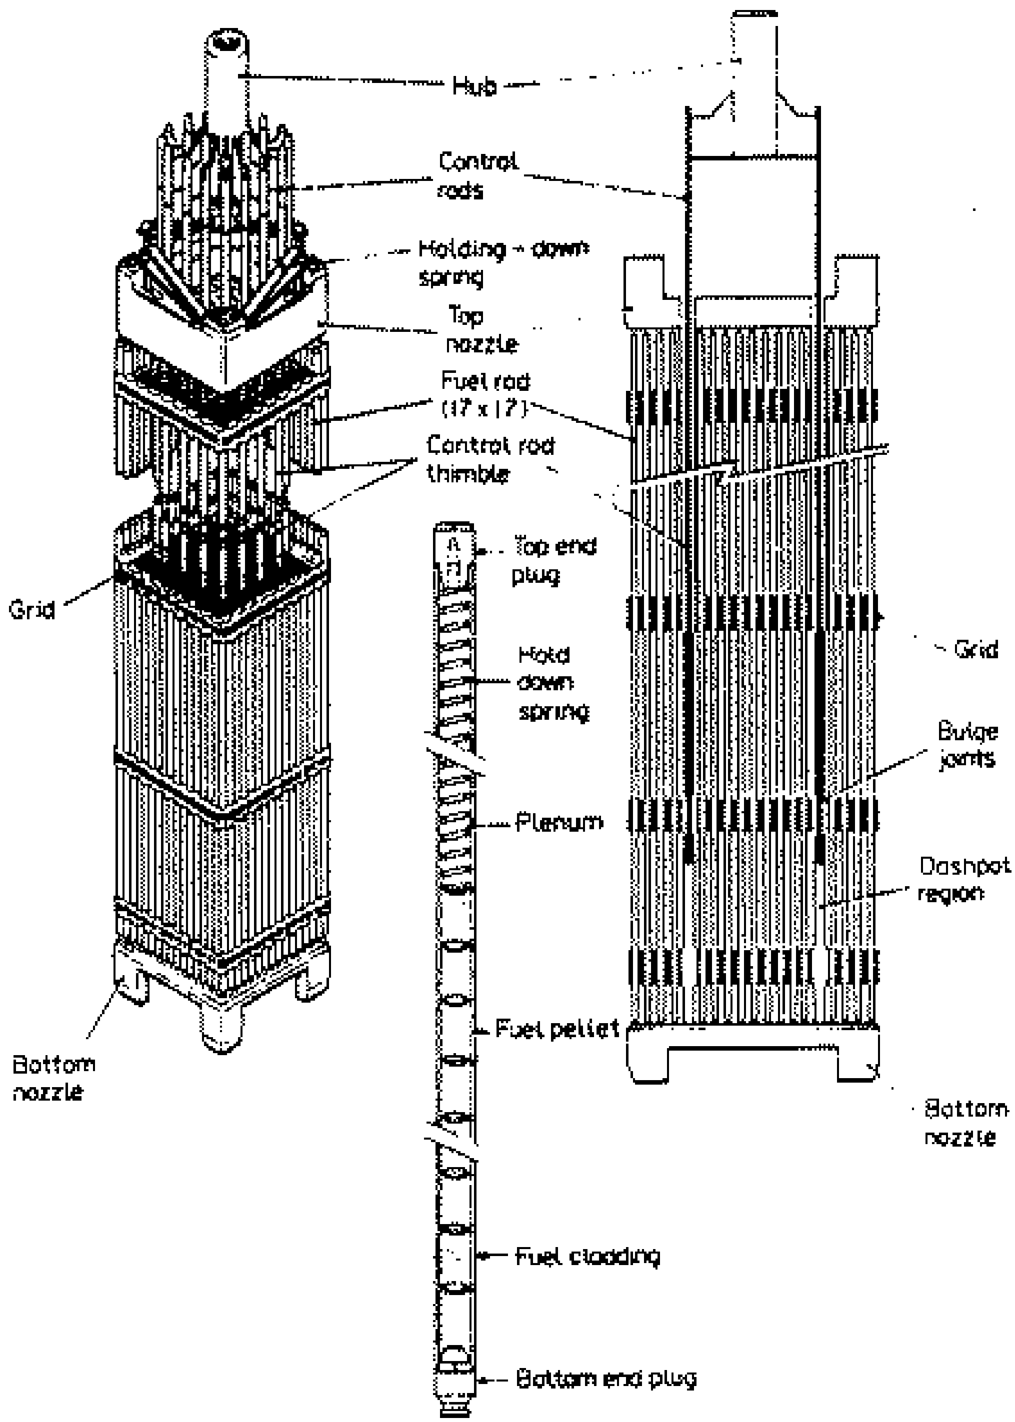

Figure 2.4 Illustration of Fuel Assemblies and Fuel Rods for a Pressurized Water Reactor

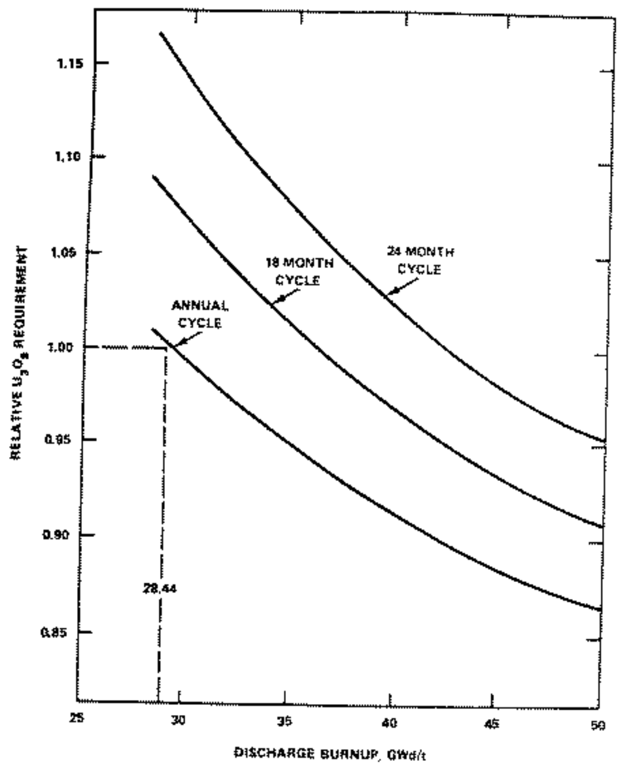

Figure 2.5 Relationship Between Planned Duration of Operating Cycle, Burnup, and Load Enrichment Levels

the NPP frequently in small batches. However this conflicts with the goal of minimizing the high opportunity cost of NPP downtime. Reactor fuel management involves a complex set of tradeoffs and the operator has many possibilities regarding the timing of refuelings, the fraction of old fuel assemblies to be replaced at each refueling, the degree of enrichment of new fuel to be added, the rearrangement of fuel within the core, and the location of moderators within the fuel assemblies. As a result, many books and articles have been devoted to this subject (see, e.g., Driscoll, Downar, and Pilat, 1990, Egan, 1984, and Silvennoinen, 1976).

The primary focus of this paper is on the optimal length of operating cycles: we assume that a decision about the planned duration of the operating spell can be implemented by an optimal fuel management strategy that maximizes fuel utilization. However the tradeoff between fuel cycle length and fuel efficiency is less important than the main tradeoff between the potential improvement in capacity utilization associated with longer operating cycles versus the potential increased risk of unplanned mid-cycle outages due to fuel failures and other other equipment failures: “The overriding consideration, however, is that the replacement of three annual cycles by two 18 month refuelings holds the possibility of reducing total outage time.” (Rahn *et al*, 1984, p. 488). Reports by the Stoller Corporation (1987 and 1989) performed an illustrative cost-benefit calculation of the gains from moving from a 12-month operating cycle to an 18-month operating cycle. Their calculations show that the improvement in capacity factor under an 18-month cycle leads to a $9.5 million reduction in the opportunity cost of lost power generation (on an annual basis). This outweighs the $6.9 million increase in annual nuclear fuel costs. However nearly half of the latter increase is due to the 500 MWh increase in output in moving from a 64.4% capacity factor under 12-month refuelings to a 68.8% capacity factor under an 18-month operating cycle. The Stoller reports concluded that the high opportunity costs of lost power generation during refueling outages favored longer 18 and 24 month refueling cycles. Further, an average NPP could increase its capacity factor by 3 to 7 percent by increasing the operating cycle from 12 to 18 months, although the gain is lower for more efficient NPPs with shorter average refueling durations. These benefits are probably the reason why "most of the U.S. utilities have adopted fuel cycle lengths longer than twelve months. Only a small number of BWRs and PWRs continue to refuel (annually)." (Stoller, 1989, p. 6-2).

Although the dominant consideration in moving to a longer operating cycle is to reduce the opportunity cost of lost power generation during a refueling, the Stoller report identified other factors that affect the utility's decision about the length of a refueling cycle:[18]

1. *Effect of NRC-mandated inspections.* Although many mandated surveillance inspections are allowed to be done on a "per refueling" basis, some inspections must be done annually or semi-annually rather than every 18 or 24 months. For example, the NRC mandates periodic in-service inspections (ISI) to assure the continuing integrity of the NPP's primary cooling system and safety-related equipment. To the extent that performing these inspections requires a mid-cycle outage, the potential benefit to longer operating cycles is reduced.[19]

2. *Preventive Maintenance Considerations.* As operating durations increase, there also could be an increase planned and unplanned outages. Although Stoller (1989, pp. 6-7) did not find evidence that rates of unplanned outages increased with the duration of an operating spell, our analysis in Sections 3 and 5 provides strong evidence that the rate of forced outages increases for operating spell durations greater than 12 months. We also show that the rate of planned outages increases with the duration of an operating spell. Both factors reduce the potential gains of moving to longer operating cycles.

[18] Besides those detailed here, the Stoller report listed several other considerations that tend to favor longer fuel cycles. These include "fewer interactions with the NRC," increased time to order reload batches, reductions in fixed maintenance and labor costs associated with each refueling, and corporate financial consideration associated with financing fuel acquisitions. EIA (1992) suggests other factors that could affect the frequency of refueling and other preventive maintenance outages: Fuel Adjustment Clauses (FACs) and utility rate of return regulations. Some states have FACs and Public Utility Commissions (PUCs) that permit rapid recovery of fuel and O&M costs through increased utility rates. Some states allow utilities to include replacement fuel costs as part of the fuel costs covered by FAC. To the extent that the FAC allows the utility to pass on the cost of refueling outages to its consumers, it might perceive a lower opportunity cost of scheduled outages. This suggests that utilities with FACs might have a lower incentive to select longer refueling cycles and have more frequent planned outages to increase the reliability of the NPP between refuelings. Paradoxically, regression results in EIA (1988) show that O&M costs for utilities facing FACs and lenient PUCs are *lower* than for utilities without FACs and facing stringent PUCs.

[19] Reactor start-ups and shutdowns are also time-consuming: Rahn *et al* (1984) estimate the average time for a cold start up of an NPP is 13 hours. A minimum of 18 hours is required for sufficient depressurization of the reactor vessel following a shutdown to enable the operator to open and inspect key components of the reactor vessel and cooling system.

3. *Manpower scheduling considerations.* Refuelings are the most labor intensive part of regular plant operations, requiring extensive pre-planning for execution of 2,000 to 3,000 separate work orders. Constraints on the availability of sufficient skilled personnel to carry out the refueling can affect the utility's decision about the frequency of refueling outages. For example, a utility with a single NPP and a work force that is only available during a single season would be more inclined toward annual refueling cycles. A multi-plant utility with summer and winter demand peaks might prefer 18-month cycles.

4. *Seasonal outage scheduling considerations.* Most utilities experience peaks in electrical demand in the summer and winter and so would prefer to schedule refuelings during the fall or spring. Eighteen month refueling cycles make it more difficult to schedule refuelings during periods of low demand whereas 12- or 24-month cycles make synchronization with low demand periods easier.

5. *Fuel management considerations.* There are many interactions between fuel cycle length and other aspects of fuel management, including "low leakage" strategies designed to reduce embrittlement of the reactor vessel. About 95% of the PWRs in the U.S. have some form of low-leakage fuel management strategy, typically involving placement of higher burnup fuel assemblies in locations next to critical welds in the reactor vessel. However the strategy of placing higher burnup assemblies in the periphery tends to reduce fuel economy and could require more frequent refuelings to maintain criticality.

6. *Radiation exposure to personnel.* Since personnel receive the highest fraction of their maximum permissible annual radiation dose during a refueling, utilities have an incentive to schedule longer operating cycles to minimize the radiation doses received by their workers.

Although nearly all utilities have explicit plans about the intended durations of their operating and refueling spells, comparisons of intended versus actual durations in Stoller (1989) reveal the prevalence of uncertain events that lead operators to deviate from their planned durations. For example, Stoller (1989, Table 1-2) shows that the actual duration of an operating cycle was generally 2 to 3 months shorter than the planned duration of the operating cycle. Figure 2.6 reproduces Figure 6-1 from Stoller (1987) which plots the distribution of the difference between actual and planned durations of refueling outages for PWRs. We see that refuelings are subject to substantial uncertainties as evidenced by the fact that 60% of refueling outages lasted 2 or more weeks longer than *ex ante* plans. This figure justifies our *exogenous refueling* specification of the DP model which posits that the duration of refueling outages is a random variable that is beyond the direct control of the operator.

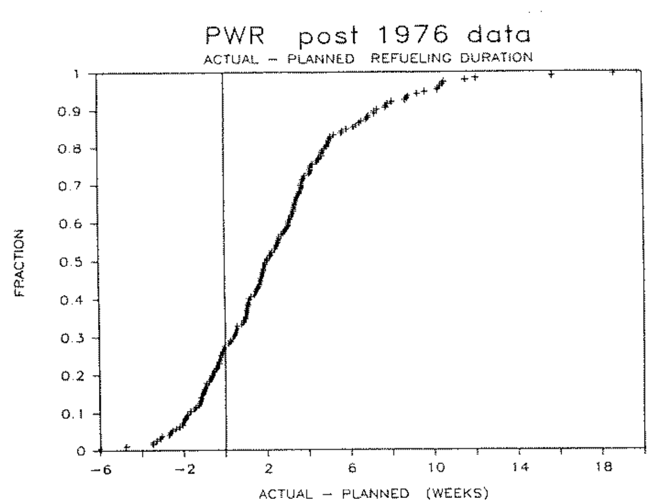

Figure 2.6 Distribution of Difference in Actual and Planned Durations of Refueling Outages for PWR's

We conclude this section by discussing longer term aging, reliability, and safety aspects of NPPs. An NPP experiences two sorts of aging problems: (1) short-term aging and exposure related-problems (e.g. fuel failures that increase with the duration of operating spells), and (2) long-term aging problems with reactor vessel, internals, steam generators, and the NPP's cooling, monitoring, and control systems. Most of the previous discussion has focused on refueling outages that succeed in partially regenerating 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 and general technological improvements in fuel reliability, instrumentation, and other aspects of nuclear power technology. The rate and duration of unplanned outages decrease monotonically with NPP age, as we will demonstrate in sections 3 and 5. 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 we have few observations of NPPs in this age range. Another problem is that while the risk of some failures is increasing with age, safety procedures are designed to make failures "rare events" which make it difficult to estimate their hazard rates from a limited number of observations.

There is a large engineering literature on problems associated with age and utilization-related deterioration in NPPs and strategies for counteracting it (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 (1992) 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 most important age-related safety hazard for PWRs and the second most important potential safety hazard for BWRs.[20] Embrittlement

[20] The primary age-related problem for BWRs was identified as corrosion in the Mark I and II containment that could lead to radioactive leakage after an accident.

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 substantial 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, discussed above.

However age related deterioration in other NPP components has forced utilities to incur huge expenses to keep their plants operational. 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 substantial downtime. Corrosion of turbine generator components is also a problem, especially for PWRs that use seawater in their secondary cooling loops. 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 lifespan. 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. These design limitations are desired upper bounds that the operator tries to avoid exceeding. 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. Many unplanned shutdowns experienced by NPPs are simply 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 would be expected to produce a bathtub-shaped pattern for age-related failures. However for reasons already discussed, it is difficult to detect the eventual upturn in the probability of age-related failures. We are able to confirm the existence of a bathtub-shaped pattern for within-cycle failures. In particular, we observe 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), and Sturm (1991). We confirm this result in our analysis in Sections 3 and 5. Simple duration analyses reveal an eventual upturn in within-cycle failure rates after 12 months. However Stoller (1989) argues that the observed upturn could be spurious, due to selectivity bias. Plants that have systematically high failure rates tend to have longer operating cycles due to the accumulation of mid-cycle outages. The higher failure rates observed for NPPs with long operating cycles might be capturing spurious duration dependence rather than true structural duration dependence. After controlling for firm-specific failure propensities by sample splitting methods, Stoller (1989) argued that within-cycle probabilities of forced outages are monotonically declining with duration since last refueling. In Section 5 we used fixed effects econometric methods to control for plant level heterogeneity and find that hazard rates for forced outages begin to increase after 12 months since last refueling.

3. Graphical Analysis of Monthly Nuclear Plant Operating Data

We extracted our data from the NRC's *Licensed Operating Reactors – Status Summary Report* (NUREG-0020), commonly known as the "Graybook" from the color of its early cover. This section describes the Graybook data, the methods we used to construct our sample, and the results of a graphical analysis highlighting its key features. This analysis provides a simple overview of NPP operations that motivates the specification of our structural econometric model in the next section.

The original dataset was constructed from an NRC tape for monthly data from January 1975 to September 1986, supplemented with data from U.S. NRC (1977).[21] We compared the data for consistency with information from the International Atomic Energy Agency (1974 and subsequent reports). The dataset was updated through February 1990 with the final tape version of the data and updated through December 1993 with the monthly diskettes that replaced the monthly publication of the Graybook.[22] Therefore, the period of our dataset is January 1975 to December 1993. These data have been used in Rothwell (1989), Rothwell (1990), David, Maude-Griffin, Rothwell, and Sturm (1991), Rothwell (1993), David, Maude-Griffin, and Rothwell (1994), and Rothwell (1995).

The Graybook data contain the following information: (1) the unit name; (2) the outage date; (3) the outage length in hours; (4) outage type (forced or scheduled); (5) the reason for the outage; and (6) the method of shut down (manual, manual scram, or automatic scram). The Graybook records eight possible reasons for outage: (1) equipment failure; (2) maintenance or test; (3) refueling; (4) regulatory restriction; (5) operator training or licensing examination; (6) administrative (including decisions to reduce output because of demand conditions); (7) operator error; and (8) other. Downtime associated for refueling (reason 3) constitutes a refueling outage. Downtime for all other reasons where the outage type is "scheduled" constitutes a planned outage for non-refueling preventive maintenance, surveillance, or equipment repair. Downtime when the outage type is "forced" constitutes an unplanned outage. The NRC defines a forced outage as any shutdown that cannot be delayed until the weekend (when power demand is lower). Because of this arbitrary cutoff, it follows that not all outages that the NRC defines as scheduled outages should be thought of a "planned outages."[23] Many of them could be outages for unexpected events that were not sufficiently serious to force the operator to shut down before the weekend. Although the Graybook records the total number of forced outages, the DP model we develop in Section 4 will only use a binary indicator of the event that one or more forced outages took place in the month. The NRC data also tell us whether a refueling outage was entered immediately after the NPP

[21] The regulatory basis for collecting this data is described in 10 CFR 1 Section 50.71, "Maintenance of records, making of reports" (January 1, 1988).

[22] The monthly diskettes were made available by Richard Hartfield of the NRC.

[23] The NRC uses "forced" and "scheduled" to describe outages. The IAEA uses "unplanned" and "planned" to describe outages. Although the definitions are slightly different, we will use the terms "forced" and "unplanned" interchangeably and "scheduled" and "planned" interchangeably, although subject to the caveat that some of the NRC-defined scheduled outages are actually unplanned outages as noted above.

was stopped with a manual or an automatic scram and whether the refueling period was immediately followed by an outage when the NPP was shut down with a scram.[24] However we have no information about unusual problems that might have been uncovered during a refueling outage.

We supplement the Graybook data with other information: the age of the plant in months since the start of commercial operation; the plant's vintage (which we define as whether the plant began commercial operation before or after the accident at TMI 2 in March 1979); the vendor of the plant's reactor (Babcock & Wilcox, Combustion Engineering, General Electric, or Westinghouse); and plant size, defined as the maximum dependable electrical generating capacity of the plant. Our dataset provides histories of 116 NPPs that were operating at some time between January 1975 and December 1993, yielding a total of 19,453 reactor-month observations.

One of our main concerns is to model stochastic processes governing NPP utilization rates. There are many measures of NPP utilization rates, the most popular of which is the *capacity factor.* It is the ratio of the electricity generated to the maximum dependable electricity that could have been generated during a specified period. Following Rothwell (1990), the capacity factor is equal to the product of the *capacity utilization rate* (the capacity factor when the plant is running) and the *service factor* (the percent of the period that the plant is running). The capacity utilization rate accounts for power reductions due to load following, end-of-cycle coastdowns, and reductions due to equipment problems. If there are no administrative shutdowns, the service factor is equal to the *availability factor*. Because there are so few administrative shutdowns, we use the terms "service factor" and "availability factor" interchangeably. We focus on the service or availability factor as the primary measure of productivity.

We begin our discussion of NPP productivity with Figure 3.0. It presents the distribution of plant size in our sample. We can see that plants are clustered into four groups: (1) small plants with capacities below 400 MWe; (2) medium-sized plants with capacities between 400 and 750 MWe; (3) large plants with capacities between 750 and 1,000 MWe; and (4) very large plants with capacities greater than 1,000 MWe. Medium-sized NPPs have done better than small or large plants: availability factors averaged 66% for the small plants, 74% for medium-sized plants, 70% for large plants, and 67% for the very large plants. There are also notable differences in plant performance for the reactor manufacturers. Combustion Engineering and Westinghouse PWRs have the highest availability factors at 73.3% and 72.6%, respectively. General Electric's BWRs and Babcock & Wilcox's PWRs have had lower average availability factors, 67.1% and 66.9%, respectively. A reason for the lower availability for Babcock & Wilcox plants

[24] Durations are measured to tenths of hours, i.e., to the nearest six minutes. If the duration began in one month and ended in a different month, we can calculate the start of the outage to the nearest six minutes. If the outage started and ended in a single month, we only know the duration's start day. Using information on start minutes, one can show that most scheduled outages began between 10 pm and 2 am and that forced outages have a uniform distribution during the day. Although we could probabilistically assign outages without start minute information, we have assumed that these outages begin at midnight or immediately after the previous outage if two outages occur on the same day. Because we focus on monthly data in this paper, this assumption should not present problems. In future work we hope to use daily data. Then our time-of-day assumption will become problematic.

is that NRC-mandated retrofits tended to be more extensive for these NPPs because of the problems discovered during and after the TMI accident. The TMI reactors were manufactured by Babcock & Wilcox.

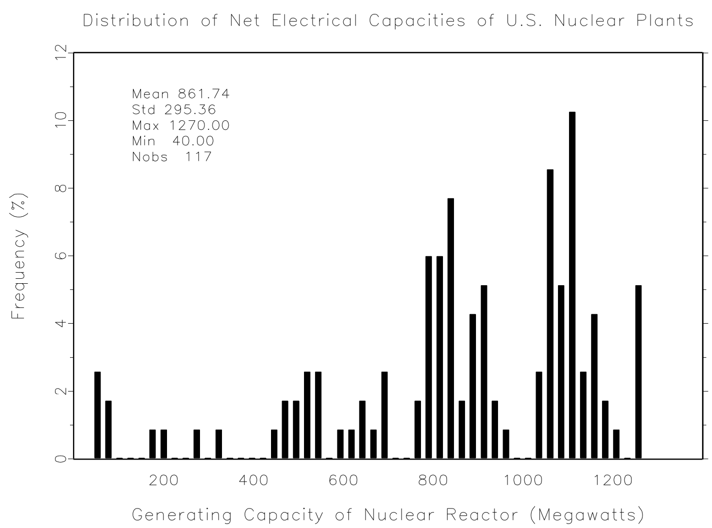

Figure 3.0 Size Distribution of U.S. Nuclear Power Plants

Figure 3.1 provides more insight into the effect of the TMI accident on NPP utilization rates by graphing the distribution of monthly availability factors over the period 1975-89, and 1980-93, respectively. The main change in these distributions is that prior to TMI the fraction of reactor/months with 0% utilization increased from about 15% before TMI to nearly 20% after TMI. It appears that utilities attempted to counteract this increased plant idleness by increasing the fraction of the time that NPP's were running at 100% availabillity, from approximately 30% before TMI to nearly 45% of the time after TMI, although when we compare the fraction of reactor/months with availability greater than 90% the change is less dramatic: 54% before TMI compared to 55% after TMI. Unfortunately the increase in the fraction of the time that NPP's were running at high availability rates was insufficient to fully offset the effect of the 5 percentage point increase in time spent at 0% availability, so that the mean availability factor fell by nearly 4 percentage points after TMI.

The pattern of changes in NPP availability factors is actually more complicated than a simple "pre/post TMI effect" as can be seen in Figure 3.2. It plots the average availability factor for the 19 years covered by our dataset. There is a sharp reduction in availability factors in 1979, the year of the TMI accident. Availability factors remained at low levels for several years following the TMI accident, but starting in 1987 availability factors began steadily increasing to the point that by 1993 NPP availability had recovered to within 2 percentage points of the 77% average availability factor in 1975.

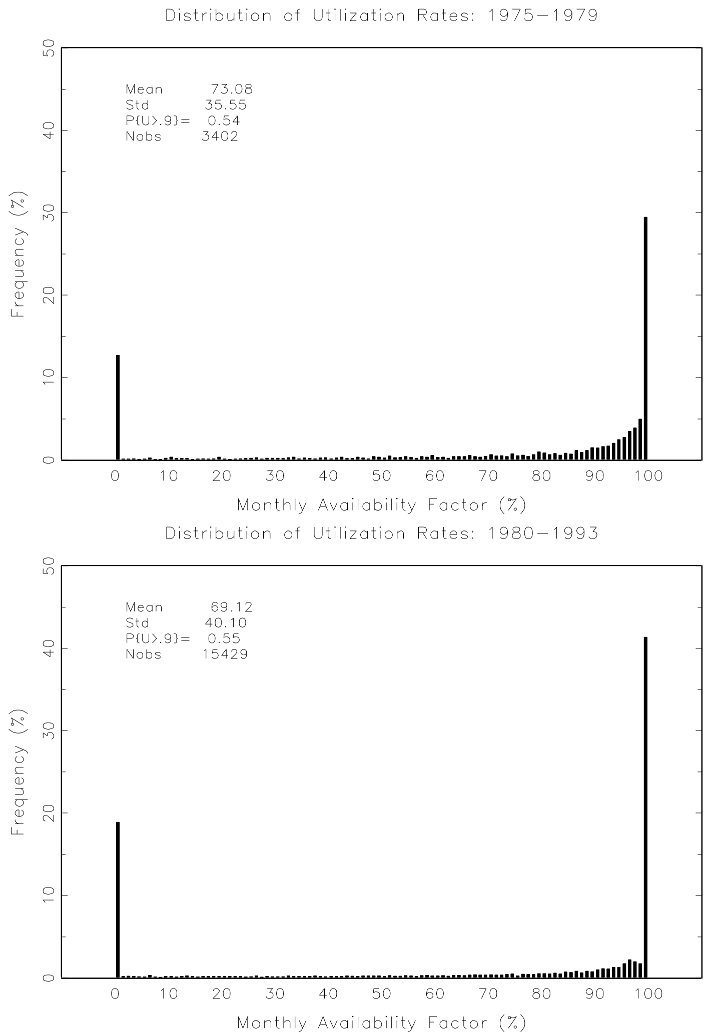

Figure 3.1 Distribution of Capacity Utilization Rates Before and After Three Mile Island

A problem with reporting average availability factors is that they hide differences in individual plant availability. Figure 3.3 illustrates the information provided by the Graybook data by plotting operating histories for the Zion-1 and Oyster Creek plants over the period of our data set, 1/75 to 12/93. We chose these two plants due to the fact that in the 1991 EIA rankings of plants based on O&M expenses, the Zion plant had the smallest level of O&M expenses and the Oyster Creek plant had the largest level of O&M expenses over the period 1974 to 1989. The vertical axes in figure 3.3 plots the monthly availability factors for these plants. The figures reveal the characteristic feature of NPP operating cycles, namely the alternating pattern of operating spells and refueling outages (marked by Rs at the bottom of the graphs). In this case we see that the duration of operating spells range from 12 to 18 months and refueling outages last anywhere from 2 months to as long as 21 months in the case of the very long refueling outage between 1983 and 1985 for the Oyster Creek plant. It is unlikely that refueling *per se* would ever take much time. Instead,

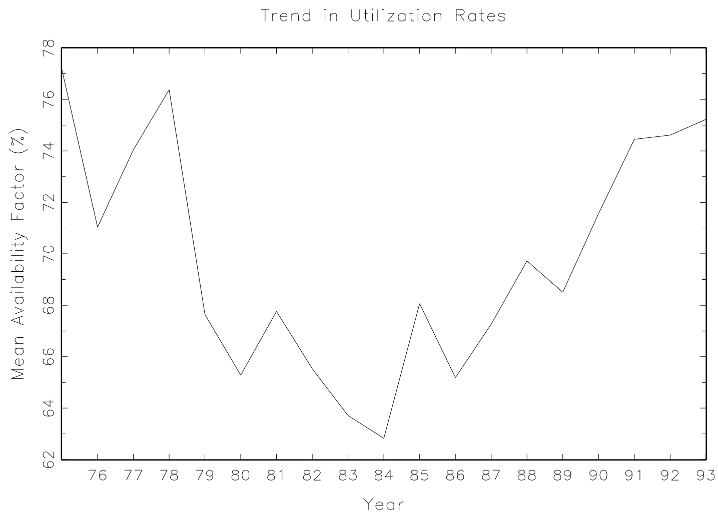

Figure 3.2 Trend in Availability Factors

these unusually long outages are examples of "major problem spells" during which major equipment repairs (such as replacement of a steam generator or installation of a mandated retrofit) occur, or capital upgrades or NRC mandated retrofits are installed. The fact that Oyster Creek experienced several major problem spells may be part of the reason for its higher O&M expenditures.

The Fs along the top of the figures denote months during which one or more forced outages occurred. As expected, capacity utilization is usually lower during these months. However we also see that there are several months where capacity utilization is nearly 100% despite the occurrence of one or more forced outages. In these cases the forced outages may have been due to "false alarms" such as the outage due to hand-held radio interference at the Duane Arnold plant mentioned at the end of section 2. Even though many forced outages may be a result of false alarms, it is perhaps disturbing to note how frequently they occur. Forced outages appear to be be a less frequent event at the Zion 1 plant than at the Oyster Creek plant, a result that may be part of the reason for Zion's lower O&M expenses. The Fs below the Rs at the bottom of the figures denote the occurrences of reactor scrams during the entry or exit from a refueling spell. As noted above, while the Graybook data give us information on reactor events leading to forced outages during operating spells, we have no information on reactor problems discovered during refueling spells other than knowing whether the reactor entered or exited the refueling spell with a reactor scram.

A final point to notice about figure 3.3 is that while majority of outages are due to refuelings and the majority of reductions in plant availability are due to forced outages, we also observe a number of scheduled outages such as the outages between 1990 and 1991 at the Zion 1 plant or the 3 month outage in early 1982 for the Oyster Creek plant. As a result of these mid-cycle outages, we observe unusually long durations between successive refuelings. For example in the case of Zion 1 one operating spell exceeded over 24 months between 1990 and 1992, and in the case

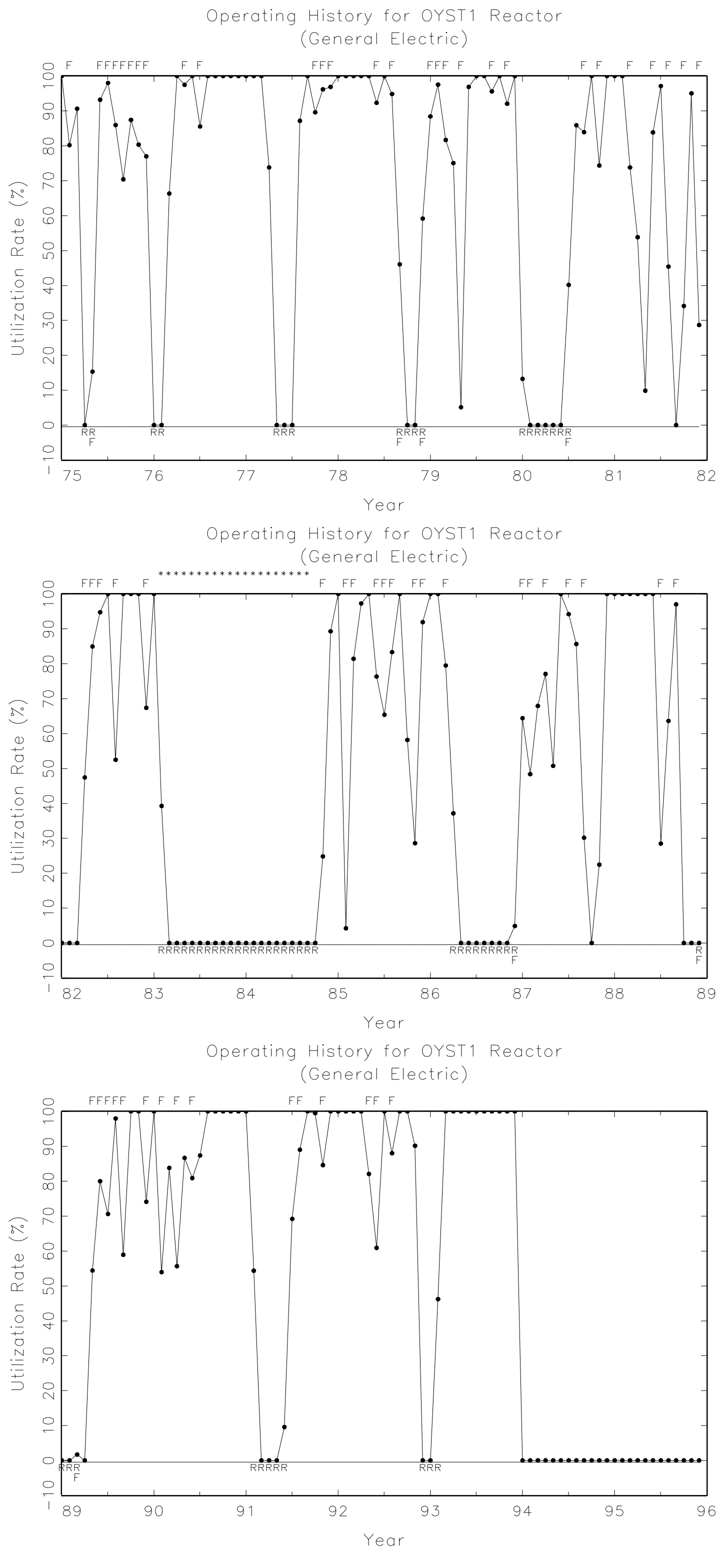

Figure 3.3 Operating History for Oyster Creek

of Oyster Creek there was an operating spell lasting more than 24 months starting in mid 1980 and ending in early 1983. There are several reasons why one should expect to observe unusually long operating cycles when there are scheduled mid-cycle outages: 1) there is no burnup of nuclear fuel during the scheduled outages so that if we measure the duration of an operating spell by the number of “effective full power days”(EFPD) between successive refuelings, the operating cycle is actually closer to 18 rather than 24 months in duration, 2) maintenance performed during the planned outage helps regenerate the plant, allowing it to continue operating longer than would otherwise be possible. This points up the importance of measuring operating spells in terms of EFPD rather than simply recording elapsed time between refuelings.

There are two main conclusions we draw from figure 3.3: a) there is considerable risk of unplanned outages resulting in large intertemporal variations in plant availability even for one of the best performing plants in the industry, b) most of the losses in NPP output are due to periodic refueling outages.

Table 3.1 confirms this by tabulating the percentage of time NPPs spent in each of four possible states that can be distinguished using the Graybook data: (1) operating; (2) down for a refueling outage; (3) down for a planned outage other than refueling; or (4) down for a forced outage. We see that refueling outages account for more than half of NPP downtime. Also, in the post-TMI period average plant availability decreased by nearly 4 percentage points. This decrease in availability was primarily the result of a 4 percentage point increase refueling outage time. Further, the fraction of time spent in planned outages decreased by 1 percentage point to 4.29% after TMI. It might initially appear from Table 3.1 that NPP operators used reductions in planned outages as the primary means of counteracting the effect of increased refueling durations. It would be misleading to conclude that the reduction in time spent in planned outages after TMI implies that utilities were doing less preventive maintenance. It is likely that the utilities were doing more preventive maintenance during the longer refueling outages. If we combine refueling and planned outages, we see that the total time spent in such outages is higher in the period after TMI: 20.76% versus 17.52%. So it is reasonable to conclude that more preventive maintenance was taking place after TMI.

	Full Sample	1975-1979	1980-1993
Percent of time operating	69.83	73.08	69.12
Percent of time refueling	15.69	12.15	16.47
Percent of time in forced outages	9.99	9.40	10.12
Percent of time in planned outage	4.48	5.37	4.29
Total Reactor/Hours	13757568	2485176	11272392
Total Reactor/Months	18831	3402	15429

Table 3.1 Summary of U.S. Nuclear Power Plant Operations

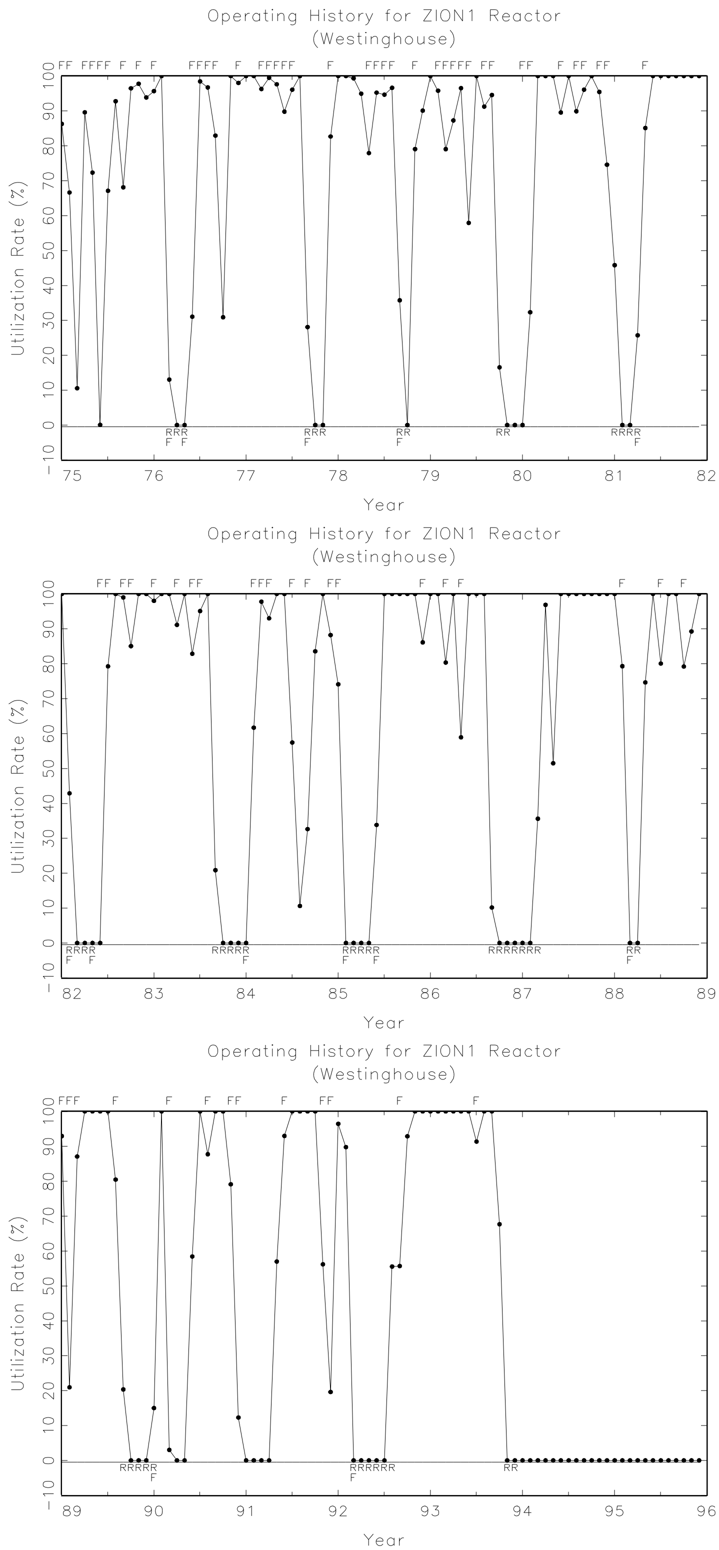

Figure 3.3 Operating History for Zion 1

Figure 3.4 gives us insight into what caused the reduction in NPP operating time following TMI. The top panel of Figure 3.4 plots the trend in average refueling durations. We see a dramatic increase in average refueling durations from a low of 1,400 hours in 1978 to a high of nearly 2,500 hours in 1983. Given that refueling is responsible for most of NPP downtime, it follows that this large increase in refueling durations is the major factor behind the decline in NPP availability after 1979. After 1983 refueling durations began a steady decline, although in 1993 the 1,800 hour average was still significantly longer than the average observed in the 1970's. The second panel of Figure 3.4 plots the average operating duration. The figure shows that there has been steady improvement in operating duration since 1979. A plot of the average operating duration, net of hours spent in outages, gives the same pattern (not shown). The improvement in the net operating duration is a combination of two factors: (1) increases in the planned length of the operating cycle from 12 to 18 months; and (2) reductions in the number and duration of mid-cycle outages.

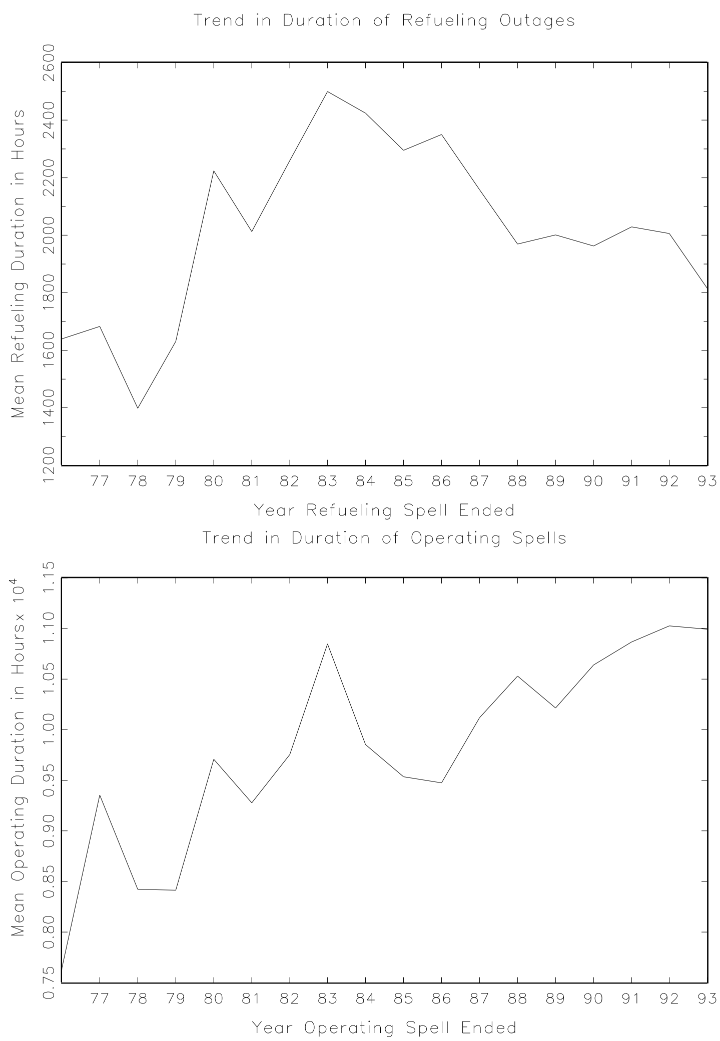

Figure 3.4 Trend in Mean Durations of Refueling and Operating Spells

Figure 3.5 confirms the industry's move toward longer operating spells by plotting the (discretized) distribution of the number of total months an NPP operates between successive refuelings. The discretized measure of operating duration does not subtract time spent in forced or planned outages. However, the discrete measure partially accounts for mid-cycle outages by not counting any month where the NPP availability is 0% due to an outage. There are some inevitable classification errors that arise from our binary discretization of months into refueling or operating spells: many refuelings begin or end midway through the month forcing us to adopt an arbitrary threshold for classifying whether a given month was a refueling or operating month. For Figure 3.5 we adopted the conservative approach of classifying any month with any hours spent in refueling as a "refueling month." So, the discretized operating spell measure necessarily understates the true operating duration. (We verified that there is essentially no change in the results from moving to a continuous-time duration analysis with operating durations measured in hours.)

There is a marked shift in the distribution of operating spells from a mode of 10 months and an average of 11.6 months through 1979 to a mode of 15 months and an average of 14.1 months after 1979. If we measure operating spells in hours (to avoid the discretization errors described above), we obtain average durations of 11.9 months through 1979 compared to 14.3 months after 1979. (This shows that the discretization errors involved in conducting the monthly analysis are negligible.) In view of the understatement of lengths of operating cycles due to discretization (and in view of our discussion in Section 2 that realized operating durations were systematically 2 to 3 months shorter than planned operating durations), it follows that Figure 3.5 is consistent with an increase in planned operating durations after 1979. Finally, there is little evidence that NPPs adopted 24-month operating cycles after 1979.

An important question is whether the move to longer operating cycles has come at the expense of a decrease in plant reliability as measured by increases in the forced outage rate. Figure 3.6 plots the trend in the average number of forced outages per month. Plant reliability was improving over time, decreasing from a rate of 0.95 forced outages per reactor-month in 1975 to approximately 0.22 forced outages per reactor-month in 1993. However Figure 3.6 does not provide a complete picture of the degree of improvement in operating reliability since it doesn't consider whether the improvement in the rate of forced outages might have been partially offset by an increase in the average duration in downtime per forced outage.

Figure 3.7 addresses this issue by plotting downtime between refuelings. Although there is still a decrease in these rates, it is less dramatic than the trend in Figure 3.6. Both panels of Figure 3.7 show pronounced increases in the time lost from outages immediately following the TMI accident: forced outage downtime hit an all time high of 14% in 1979 and planned outages hit an all time high of 8% in the following year. The decrease in the time lost due to planned outages is greater than the decrease in downtime due to forced outages. Planned outages consumed approximately 6% of mid-cycle availability in 1975 and decreased to only 2% of mid-cycle unavailability by 1993. Forced outages consumed about 9.5% of mid-cycle availability, decreasing by only a quarter of a percentage point to

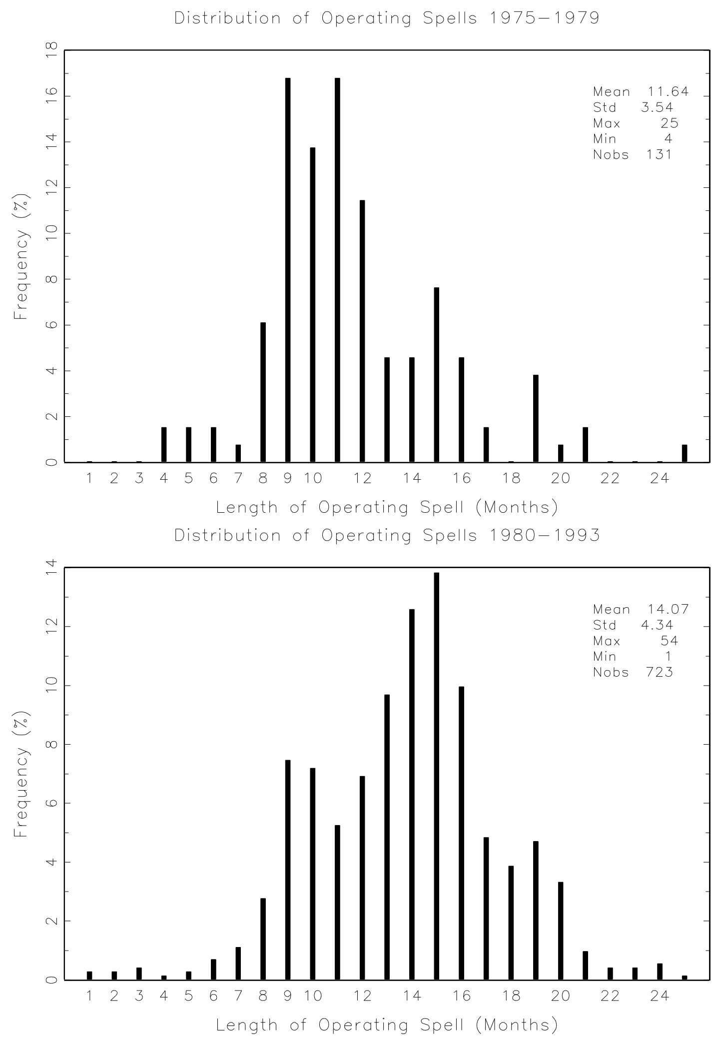

Figure 3.5 Discretized Distributions of Operating Spell Durations before and after TMI

9.25% in 1993. However while the improvement in lost time from planned mid-cycle outages is significant, the bulk of the improvement in average availability factors is due to the increased length of operating spells.

We now examine the issue of whether the probability that an NPP experiences one or more forced outages during a month has the characteristic "bathtub" shape (as a function of duration since last refueling) that has been discussed in the engineering literature. The Stoller reports (1987 and 1989) found that this probability is monotonically decreasing with duration since last refueling and concluded, "Clearly, longer cycles have not lead to poorer performance between refuelings." (Stoller, 1989, p. 5.2). Figure 3.8 plots the probability and the average number of forced outages during a month as a function of the operating duration. Although there is month-to-month variation in these rates, the figures provide some support for the "bathtub hypothesis." The forced outage rate is particularly high in the first month

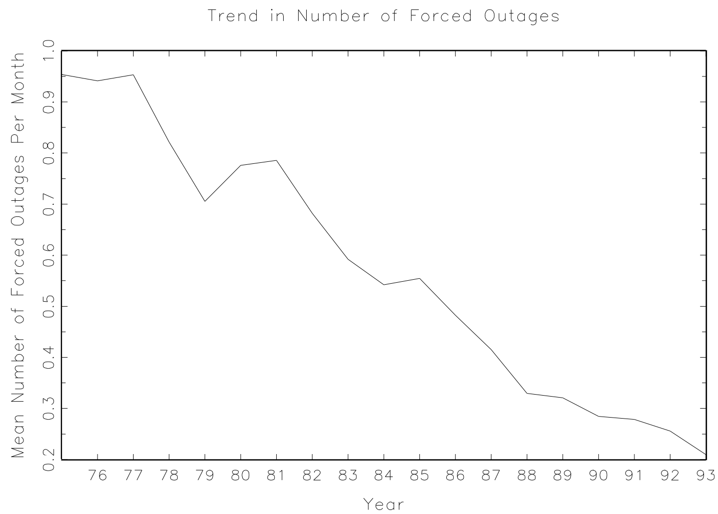

Figure 3.6 Trend in Mean Number of Forced Outages per Reactor-Month

after refueling, reflecting problems from a cold start after a refueling outage. Forced outage rates drop significantly over the next several months and appear to bottom out between the 12^{th} to the 14^{th} month of the operation and then begin increasing. There is a large variation in the estimated forced outage rate for operating spells beyond 18 months, because the number of observations drop precipitously after this point. We have only 130 reactor-month observations for operating spells greater than 20 months compared to 780 observations for the first month of operation. Despite these problems, one cannot reject the hypothesis of an increase in the forced outage rate after 14 months of operation through casual empirical observation.

Further, simply plotting the forced outage rate does not give us the full picture of the possible problems that NPPs could be experiencing toward the end of a long operating spell. These outages could be more severe, requiring outages of longer duration. In addition, the time spent in planned outages could increase. Figure 3.9 confirms this by plotting average availability factors and the average fraction of time spent in outages against the operating duration. The average availability is only 84% in the first month following a refueling. Availability then steadily increases to a maximum of 90% in the 5^{th} month of the operating spell and then begins a steady decline. There is an especially precipitous fall in NPP availability beyond 18 months of operation. The upturn that occurs in months 24 and 25 could be spurious: a result of the problems of having very few observations in this range. The second and third panels in Figure 3.9 show the bathtub-shaped loss of time for outages. These figures seem to confirm the hypothesis that forced outages that occur later in the operating cycle are more “severe” in the sense that they entail significantly greater downtime than forced outages that occur earlier in the operating spell. However one must interpret these results extremely cautiously since they are based on very few observations.

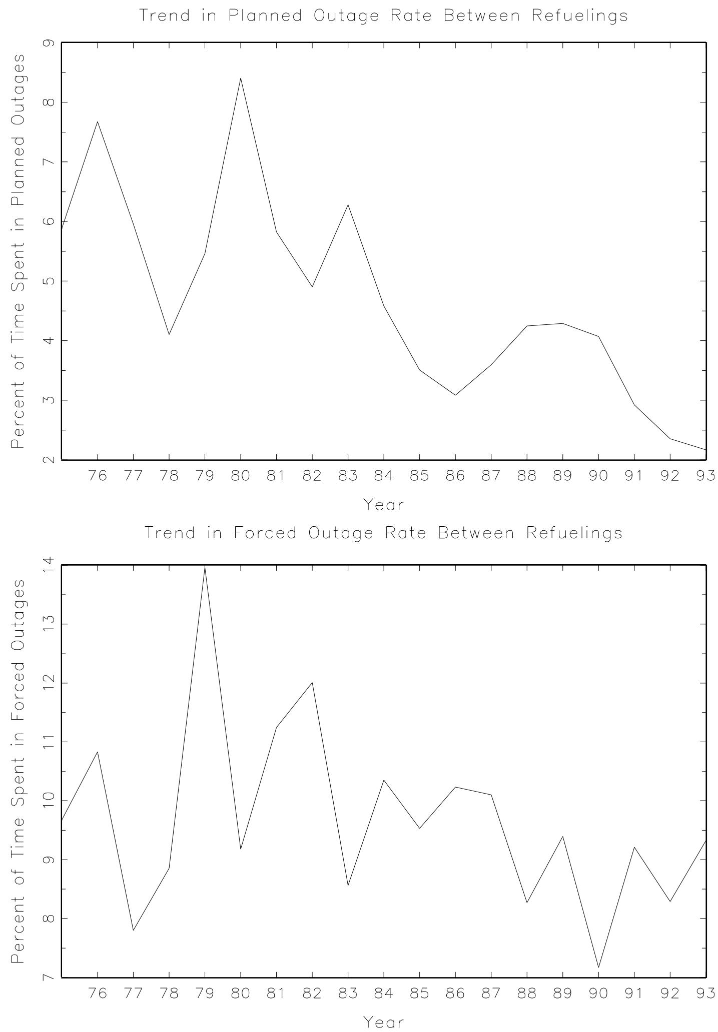

Figure 3.7 Trend in Time Lost Due to Planned and Forced Outages Between Refuelings

The Stoller reports noted that besides the problems of small numbers of observations, the upturns in the failure rate could be an artifact of selection bias described in Section 2, i.e., most of the observations on very long fuel cycles will correspond to poorly performing NPPs that experienced many mid-cycle outages. We do not think that our results reflect such selectivity bias given the way we measured the operating duration. We did not increment the operating duration in any month that an NPP had no output. Thus our measure of duration is an approximation of the number of effective full power days (EFPD) since last refueling. Our procedure could lead to an opposite form of selectivity bias, i.e., most of the observations of very long operating spells could correspond to more reliable NPPs because it would not be sensible to run an unreliable NPP on very long operating cycles.

Figure 3.8 Effect of operating durations on the monthly probability and average forced outages

We now turn to the issue of identifying the effect of NPP age on reliability. Recall from Section 2 that previous analyses have generally found systematic improvement in reliability with plant age, an effect that is usually interpreted as the result of "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 3.10 plots the monthly probability and average number of forced outages as a function of plant age. Due to the nonstationarities connected with the TMI incident, Figure 3.11 (and Figure 3.12 below) are based on observations after 1980. 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. Both figures show 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

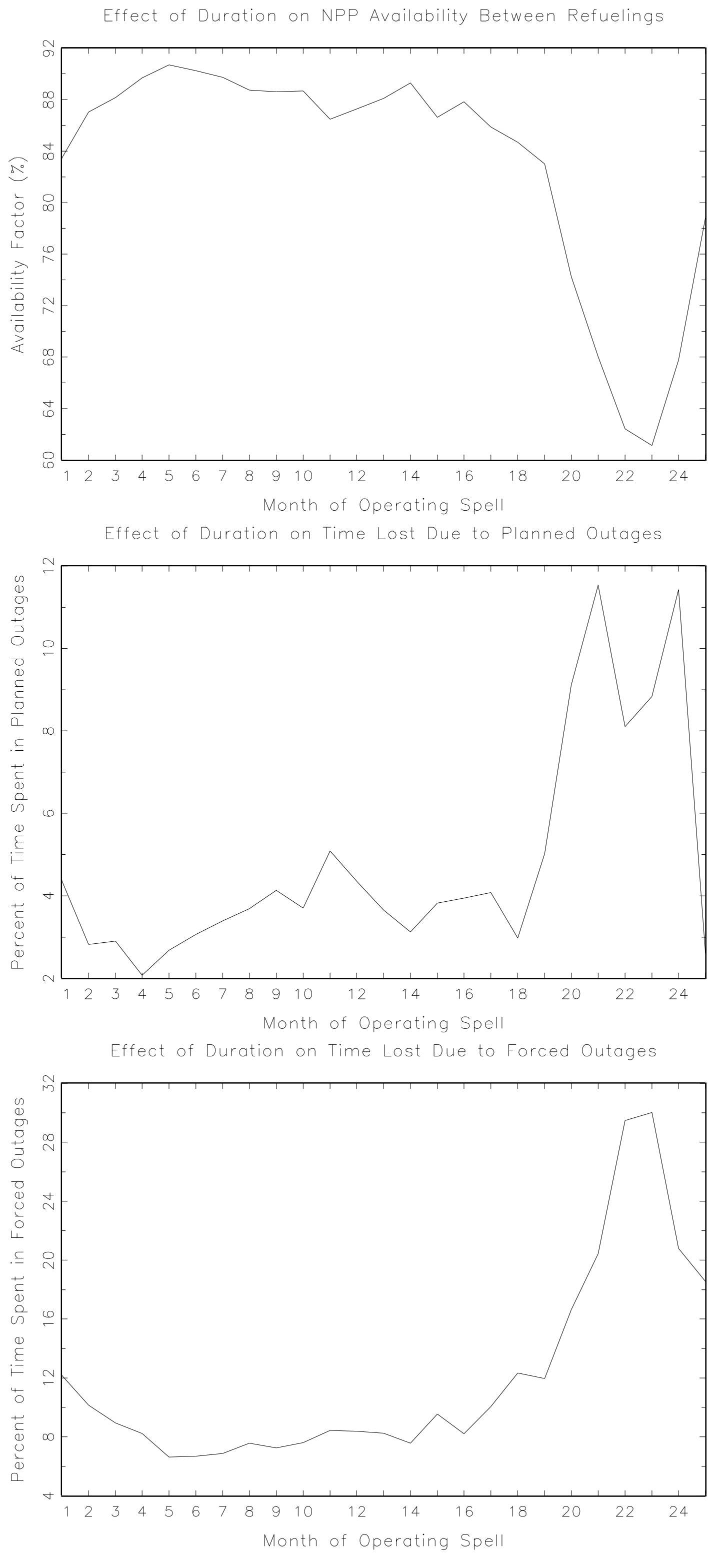

Figure 3.9 Effect of operating durations on NPP availability and downtime

having so few observations on NPPs over 24 years old. Indeed, in our sample there are only 4 NPPs that were over 25 years old. Thus one needs to interpret the estimates of the tails of these age-related probabilities with extreme caution.

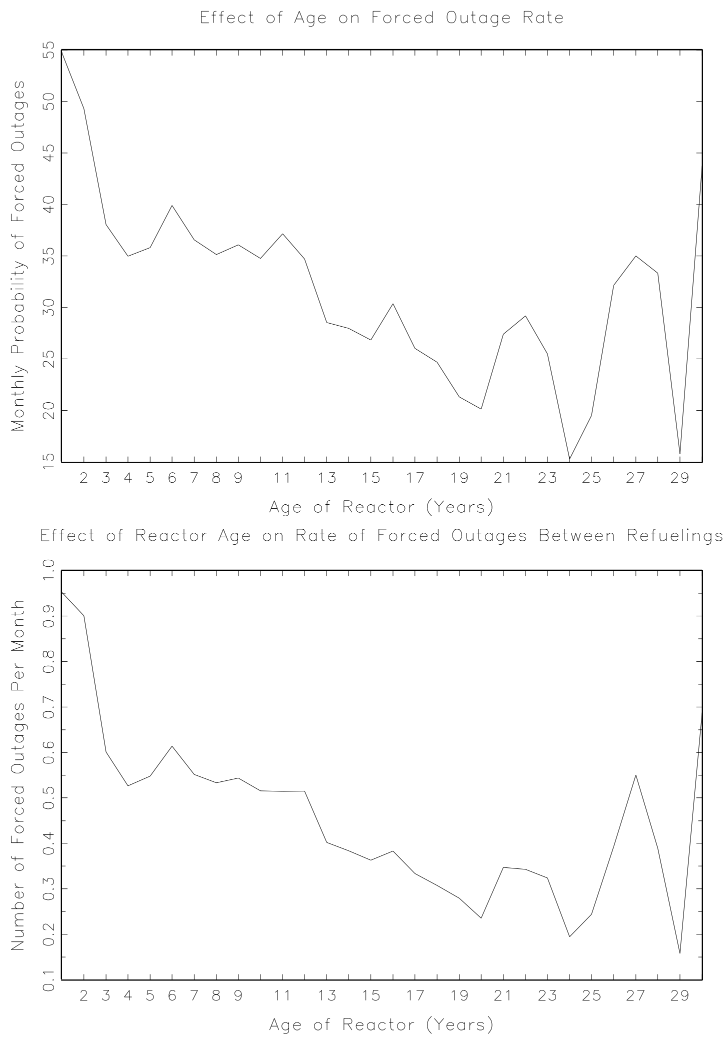

Figure 3.10 Effect of Reactor Age on the Probability and Mean Number of Forced Outages after 1980

Figure 3.11 plots the relationship between NPP availability factors (excluding refuelings) and the fraction of downtime (excluding refuelings). We see a marked improvement in the availability factor after the first year of operation after which performance is independent of age. The lower two panels of Figure 3.11 show that after the large declines in the fraction of time spent in outages after the first year of operation, subsequent improvements are much more gradual. Again, we do not put much weight on the estimates of these curves beyond 25 years of operation.

When we replot Figure 3.11 to include all months, rather than only operating months between refuelings (not shown), we find that the curve has a bathtub shape, decreasing for the first 15 years of operation and then increasing. Plots of the fraction of time spent in outages show much less of a tendency to decline with age when we compute them over all months as opposed to only those months between refuelings. In particular, we see large peaks in the percent of time spent in forced outages at 6 and 14 years of operation and similar peaks in the percent of time spent in planned outages at 10 and 16 years of operation. These peaks could represent time spent in long outages for major equipment repairs.

Further, Figure 3.12 shows the effect of age on the operating durations and refueling outages. We see significant increases in the operating spells after the second year of operation and gradual increases thereafter. There is some evidence of a tendency for declining durations after 19 or 20 years, although we make this inference cautiously given the lack of observations on NPPs over 20 years old. The second panel shows that there is no effect of age on the duration of refueling outages. This is surprising given that one would expect more preventive maintenance and repair activities for older NPPs, leading to longer refueling durations. On the other hand, calendar time is highly collinear with NPP age. Many age effects could be confounded with time-related improvements. In particular, the steady decrease in the forced outage rate could reflect the effects of improvements in fuel reliability over the 1980s.

In summary, our simple graphical analysis illuminates key aspects of NPP operation. We have found that most of the downtime is spent in periodic refueling outages. However, all NPPs face a significant risk of forced outage. These outages have a 35% probability of occurring in any given month and consume 10% of the potential time an NPP could be operating. We have documented large changes in NPP operation over time. A big shift in operating behavior occurred in the 1980s. The main changes have been (1) an increase in the duration of refueling outages; (2) a shift in the planned duration of operating cycles from 12 to 18 months; and (3) a steady decrease in the forced outage rate.

There is evidence that the probability an NPP experiences forced outages has a classic "bathtub shape" with the monthly probability of experiencing one or more forced outages starting at nearly 50% in the first month following a refueling outage, decreasing until about the 12^{th} month of an operating spell, and then increasing. We have also found some weak evidence that outages that occur later in the operating cycle may involve more downtime, possibly reflecting an increasing share of serious problems as opposed to false alarms. However we put very little weight on this finding since our estimates are based on very few observations. We are unable to find strong evidence for a bathtub-shaped pattern for forced outages as a function of NPP age. Although we see large increases in the reliability of the NPP after the first several years of operation, the more gradual improvements in NPP reliability could be a result of confounding influence of technological progress, such as improvements in instrumentation and fuel reliability. Given the small number of observations on older NPPs, our simple graphical analysis of the Graybook data was unable to detect any evidence of systematic age-related deterioration in NPP performance.

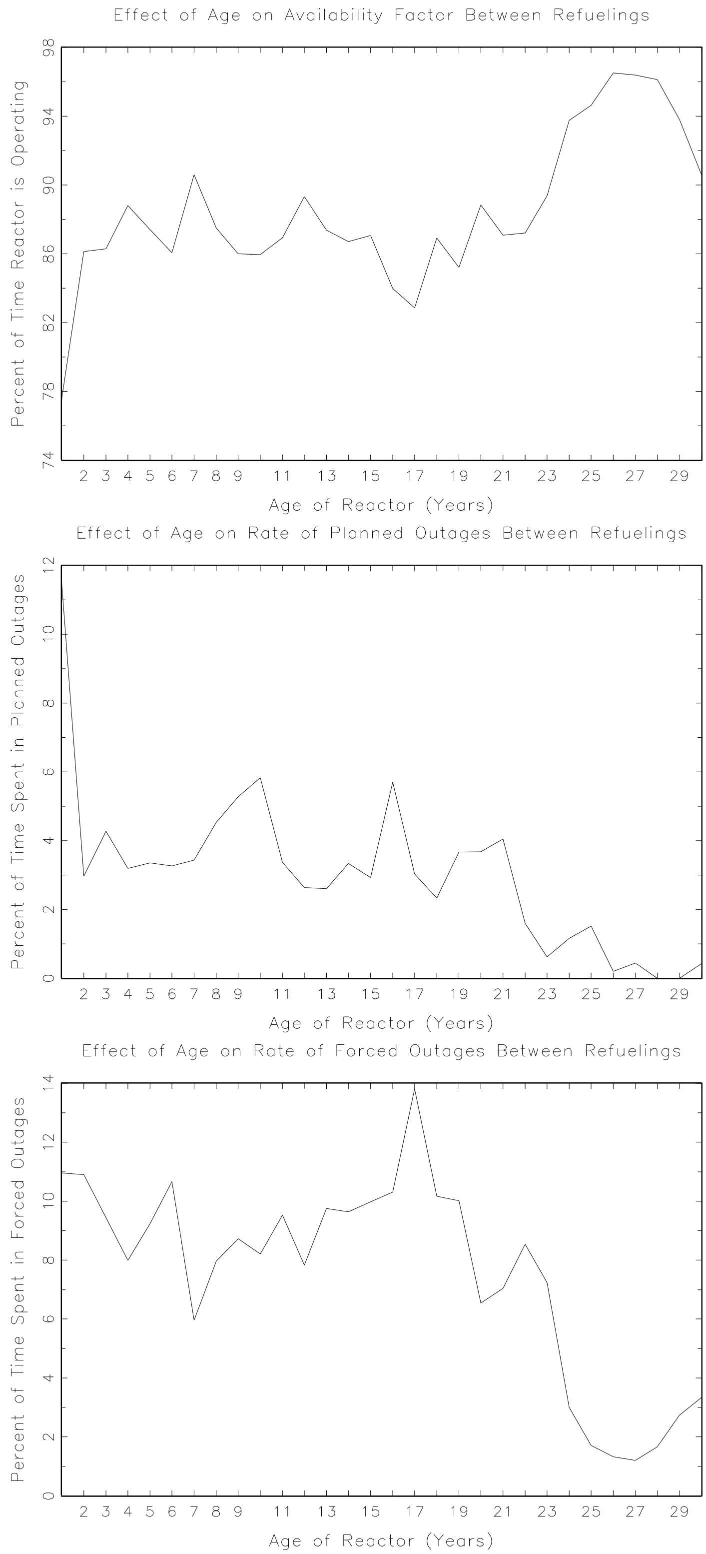

Figure 3.11 Effect of Reactor Age on NPP Availability and Downtime after 1980

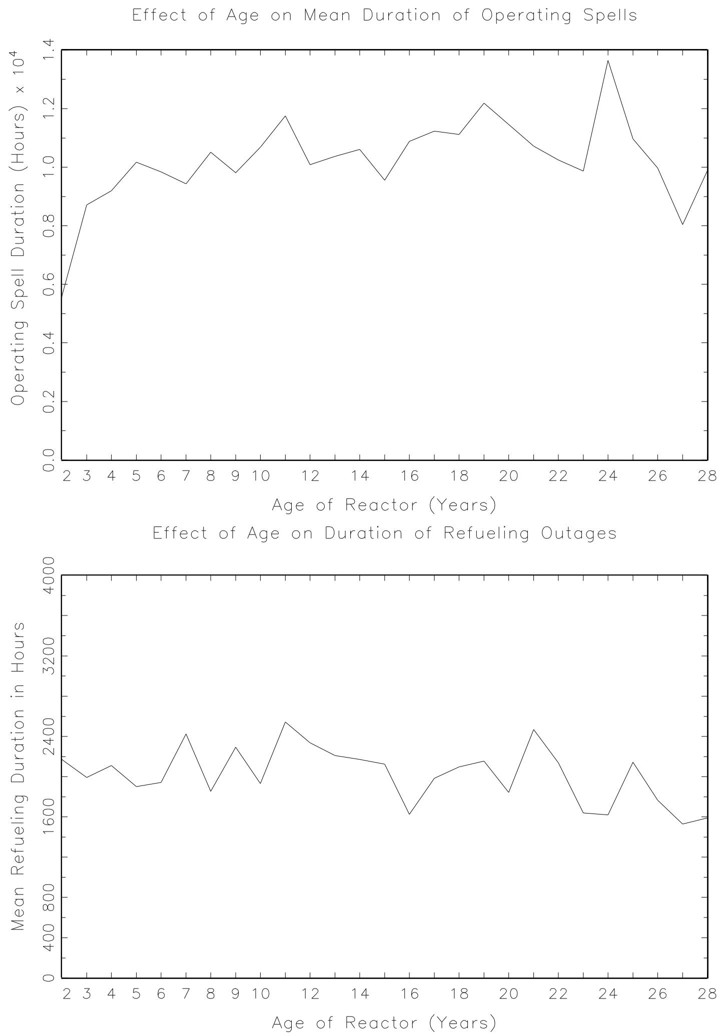

Figure 3.12 Effect of Reactor Age on Durations of Operating Spells and Refueling Outages

4. A Dynamic Programming Model of NPP Operations

In this section we formulate a dynamic programming (DP) model designed to capture the main features of NPP operations discussed in Section 3. 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 constraints and regulatory constraints imposed by the NRC and state public utility commissions. The DP specification consists of a vector of *state variables,* s_t, a *control variable,* a_t, a *profit function* $\pi(a, s)$, a *discount factor* $\beta \in [0, 1]$, and a transition density $\lambda(s'|s, a)$, representing the stochastic law of motion for the state of the NPP.

The DP problem is in discrete time and NPP operating decisions are assumed to be made at the start of each month. We assumed monthly intervals so the predictions of the DP model match 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 such as adjustment of control rods and concentrations of moderators in the reactor coolant. We also abstract from the complicated details of the fuel management strategy discussed in Section 2. Instead, our model focuses on the larger scale decisions about plant utilization and shutdowns, the timings of planned outages for refuelings and preventive maintenance, and whether a plant should be closed for decommissioning. We assume that any strategy specifying the "big picture" decisions governing NPP operations can be implemented by an optimal continuous time strategy that specifies the minute-by-minute details. These include controlling the reactivity of the core, the specific repair and inspection activities during forced outages, and the fuel management strategy during refueling outages. Although day-to-day management involves complex tradeoffs, given the high opportunity costs of NPP downtime, these complexities are secondary in importance to the larger issue of determining the optimal length of the fuel cycle. Our DP model is designed to capture the most important tradeoffs that plant operators face in making the primary decisions affecting NPP operations.

A key advantage of the DP framework is that it allows us to determine optimal utilization strategies that account for the impact of uncertain events such as equipment failures and other signals indicating abnormal conditions in the reactor and generator. Since these signals affect the operator's utilization decision in continuous time, we interpret the monthly utilization decision in our discrete time DP model as corresponding to the integral of the instantaneous utilization decisions over the current month. Similarly we interpret the state variable as reflecting the entire stream of signals that the NPP operator receives during the month. We assume that the operator has perfect control over the level of productivity but imperfect control over equipment failures and the NPP's abnormal status signals.[25] However

[25] An automatic scram results in a shutdown that is typically beyond the operator's control. However unless the scram was as result of a false alarm or equipment failure, it is likely that the operator would have voluntarily made the same decision (known as a manual scram). Even if an automatic scram was initiated by a false alarm, the operator has the option of rapidly resuming operation from a "hot start" rather than keeping the NPP shut down. Therefore, our assumption that the operator has perfect control over the reactor seems reasonable, even though there have been rare occasions, such as the accident at Three Mile Island, where the reactor has been "out of control."

through equipment repair and preventive maintenance, the operator can partially "regenerate" the NPP, presumably reducing the probability of receiving abnormal status signals and of experiencing equipment failures in the future.[26]

The operator clearly observes more signals about the plant'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 the econometrician only observes x_t. We can think of the unobserved state vector as reflecting the myriad of information displayed in the NPP's control room. This information includes the reactor's temperature, pressure and neutron flux; status checks on hydraulic valves and pressure relief devices; measurements of coolant flow, water chemistry, and radioactivity levels; electrical output and transients; etc. 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, so that $\epsilon_t(a)$ represent the operator's assessment of the net cost or benefit to taking action a conditional on all available information. Since the full set of information available to the NPP operator is unobserved, we treat ϵ_t as a latent random vector with a distribution we will specify below.

We follow the general framework of Rust (1988) and assume that the operator's current period profit from taking action a for an NPP, which is in current state (x, ϵ), is given by the function $\pi(a, x, \epsilon)$. It has the additively separable representation

$$\pi(a, x, \epsilon) = \mu(a, x, \phi) + \epsilon(a), \tag{4.1}$$

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 plant operator chooses an optimal operating strategy $a_t = \alpha_t(x_t, \epsilon_t)$ that maximizes the plant's expected net present value $V_0(x, \epsilon)$ given by

$$V_0(x, \epsilon) = \max_{(\alpha_0, \ldots, \alpha_T)} E\left\{ \sum_{t=0}^{T} \beta^t \pi(a_t, x_t, \epsilon_t) \middle| x_0 = x, \epsilon_0 = \epsilon \right\}. \tag{4.2}$$

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. Although there is a possibility that the operator could apply to the NRC for a

[26] This is an oversimplification in view of the bathtub-shaped pattern of forced outage probabilities described in Sections 2 and 3. In particular, there is an unusually high rate of unplanned outages in the first month following the cold start of an NPP, but this forced outage rate decreases until 12 months into the operating cycle when it begins increasing again. Our numerical solutions to the DP model fully account for complexities induced by the bathtub-shaped pattern of forced outage probabilities.

license extension, no plant has been granted an extension, so we have no observations of this in our sample. Therefore, we will initially assume a 40-year life, which corresponds to $T = 480$ in our monthly DP model.[27]

We adopt an important simplifying assumption that the transition density λ can be factored as

$$\lambda(x_{t+1}, \epsilon_{t+1}|x_t, \epsilon_t, a_t) = p(x_{t+1}|x_t, a_t, \psi)q(\epsilon_{t+1}), \tag{4.3}$$

where ψ is a vector of unknown parameters characterizing the transition density for the observable part of the state and control variables. Equation (4.3) 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:

$$\begin{aligned} P_t(a|x) &= \int I\left\{a = \alpha_t(x, \epsilon)\right\} q(d\epsilon) \\ &= \frac{\exp\{v_t(x,a)\}}{\sum_{a' \in A_t(x)} \exp\{v_t(x,a')\}} \end{aligned} \tag{4.4}$$

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

$$v_t(x, a) = \mu(x, a, \phi) + \beta \int \log\left[\sum_{a' \in A_t(x')} \exp\{v_{t+1}(x', a')\}\right] p(dx'\Big|x, a, \psi). \tag{4.5}$$

The set $A_t(x)$ represents the set of feasible actions available to the operator in state x at time t and will be specified shortly.

We compute the solution to the DP model using standard backward induction methods with the recursion equation (4.5). The implied stochastic process for the observed state and control variables $\{x_t, a_t\}$ constitutes the DP model's prediction of the optimal strategy for running, refueling, and closing an NPP. These predictions, however, depend on a vector $\theta = (\beta, \phi, \psi)$ of unknown parameters specifying the discount factor, the unknown parameters of the profit function, and 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 a sample of 114 U.S. NPPs. Denote this data by $\{x_t^i, a_t^i\}$, $t = \underline{t}_i, \ldots, \overline{t}_i$, $i = 1, \ldots, 114$. Given the conditional choice probability $P_t(a|x)$ in equation (4.4) and the decomposition of the transition density $\lambda(x', \epsilon'|x, \epsilon, a)$ in equation (4.3), it is straightforward to estimate the unknown parameter vector $\theta = (\beta, \phi, \psi)$ by maximum likelihood using the (full) likelihood function

$$L_f(\theta) = \sum_{i=1}^{114} \sum_{t=\underline{t}_i+1}^{\overline{t}_i} \log\left[P(a_t^i|x_t^i, \theta)p(x_t^i|x_{t-1}^i, a_{t-1}^i, \psi)\right]. \tag{4.6}$$

[27] For similar reasons we have ruled out the possibility that the operator will undertake major investments, such as thermal annealing, designed to extend the life of the NPP. We are not aware of such investments to date. Given that there is much uncertainty about the potential benefits of these actions (which make it less likely that these investments would be profitable), we feel it is reasonable to ignore them.

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

$$L_1(\psi) = \sum_{i=1}^{114} \sum_{t=\underline{t}_i+1}^{\overline{t}_i} \log \left[p(x_t^i | x_{t-1}^i, a_{t-1}^i, \psi) \right], \tag{4.7}$$

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

$$L_2(\beta, \phi|\hat{\psi}) = \sum_{i=1}^{114} \sum_{t=\underline{t}_i+1}^{\overline{t}_i} \log \left[P(a_t^i | x_t^i, \beta, \phi, \hat{\psi}) \right]. \tag{4.8}$$

Rust (1988) established the consistency and asymptotic normality of the two-stage and full maximum likelihood estimators.[28]

In the remainder of this paper we will adopt a structural approach to inference and interpret observed NPP operating histories as realizations of controlled stochastic processes. The reduced-form analysis in the previous section provided strong evidence that the stochastic processes governing plant availability and forced outage rates are non-stationary. In particular, we showed that there is clear evidence of duration, age, and time dependence. The DP model can provide simple explanations for observed patterns of duration and age dependence, but it is more difficult to develop a model that accounts for variations in operating behavior with calendar time. Modelling time dependence requires specifying the stochastic processes governing regulatory stringency, electricity prices, and technological progress. We have not attempted to do this in this paper. However, the analysis in Section 2 suggests that the primary effect of calendar time on NPP operations is associated with the shift to a stricter regulatory regime following the TMI accident in 1979. We account for this by restricting our sample to the period 1980-1993. In future work we plan to use this model to predict the effects of alternative regulatory policies on the operating behavior and discounted profits of NPPs.

This paper concentrates on showing that our DP approach provides a credible model of NPP operations in the post-TMI era. Although we would like our model to predict the impact of technological changes, as we noted in the previous section, it is extremely difficult to disentangle age and time effects in our dataset. Thus we have not attempted to identify the separate effects of technological progress, plant level "learning by doing," efficiency improvements associated with transition to an equilibrium fuel cycle, and age related deterioration. The DP model does allow us to separate the combined effects of these factors from the "horizon effect" created by the 40-year NRC operating license. We caution the reader that any long-term forecasts of electricity output from this version of the DP model are based on the implicit assumption that there will be no major changes in electricity prices, technology, or regulatory policies throughout the forecasting period.

---

[28] The covariance matrix for the parameters (β, ϕ) will not be consistently estimated from the second stage partial likelihood due to estimation noise in the first stage parameters $\hat{\psi}$. One can use 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 with fully efficient estimates of θ.

Another issue is how to deal with plant-level heterogeneity. Despite the evidence of systematic plant-specific differences in NPP performance and reliability given in David, Maude-Griffin, and Rothwell (1994), it is difficult to find observable factors that correlate well with these differences. This suggests that the returns to developing a model that accounts for heterogeneity could be small, especially in view of the substantial econometric and computational problems involved.[29] Therefore, we have decided to start with a simple model that treats all NPPs as homogeneous. In future work we plan to generalize our model to account for effects such as plant size, reactor manufacturer, and differences in local regulatory climate.

Our DP model can be viewed as an extension of the model developed in the recent paper by Sturm (1994) that was applied to the operation of European NPPs. Sturm's paper focused on modelling operating spells in a framework that treats observed operating histories of NPPs as censored point processes. In Sturm's discrete time model an NPP can be in one of three possible states in each period: (1) operating; (2) shut down due to forced outage; or (3) shut down due to refueling outage. (Also see Rothwell, 1989, for a similar three state analysis.) If the NPP is currently operating the plant operator has two possible decisions: continue operating or shut the plant down for refueling. If the plant is down due to a forced outage, the plant operator has two possible decisions: bring the plant back online or begin refueling. Sturm modelled the refueling duration as an exogenous (geometrically distributed) random variable: "The combined effects of the plant manufacturer's recommendations, 'sound engineering practice', and regulatory constraints may be such that the refuel duration is a random variable beyond the immediate control of the operator." (Sturm, 1994, p. 10-11).

Our model differs from Sturm's model in the following ways: (1) we use a finite horizon model, whereas Sturm used an infinite horizon model; (2) we use a three-dimensional state variable, x_t, whereas Sturm's model had a single state variable, x_t, indicating whether the plant was in one of the 3 states listed above; (3) we model the operator having 7 possible utilization/refueling decisions as opposed to the binary operate/don't operate decision in Sturm's model; and (4) our model gives the plant operator the option of closing an unprofitable NPP for decommissioning.

In our DP model there are two ways of interpreting refueling outage durations. In the *exogenous refueling model* refueling durations are predetermined random variables whose realizations are beyond the control of the operator for the same technological and regulatory reasons as in Sturm's model. In the *endogenous refueling model* refueling durations are completely within the control of the operator. In this model the refueling duration is interpreted as an endogenously determined random variable, a result of a cost-benefit calculation by the NPP operator. In both the exogenous and endogenous models the decision to begin refueling is made by the operator. The difference in the two treatments is the issue of how best to model durations of refueling spells once they are initiated. In the exogenous

[29] Accounting for heterogeneity essentially involves separate solutions of the DP problem for each different "type" of NPP. Because of the unique characteristics of each NPP, each could be considered a different type. See David and Rothwell (1995) on the lack of standardization in the U.S. nuclear power industry as compared to the French nuclear power indusrty.

refueling model the operator has only one decision during refueling: whether to close the plant. If the operator decides not to close it, refueling must end before the plant output level can be determined.

Without good information on technological and regulatory constraints influencing the refueling duration, it is difficult for the DP model to explain why an operator would ever choose to be down for refueling for more than the minimal time, given the high opportunity costs of NPP downtime. As mentioned earlier, the Graybook data do not provide information on potential problems discovered in surveillance inspections during refueling outages, information that could be useful in explaining why most NPPs have refueling outages that last significantly longer than the minimal time needed for recharging the core. Without such information the DP model will attempt to fit the data by making refueling appear to be a "profitable" rather than a very costly activity. In addition, the Graybook data do not provide information on burnups of discharged fuel elements or the amount of new fuel loaded at each refueling. So, we have no direct way of measuring fuel costs and fuel efficiency other than the duration of the previous operating spell. Therefore, it is difficult to capture the tradeoff between fuel costs and the length of the operating cycle. Our DP model does allow us to capture what we view as the fundamental tradeoff: longer operating cycles enable the operator to reduce the opportunity costs of plant downtime but at the cost of an increased risk of forced outages associated with cycle-specific deterioration in the reactor and fuel elements.

We now turn to a detailed description of the (observed) state and control variables used in our model.

State Variable $x_t = (r_t, f_t, d_t)$ 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;

$f_t = 3$ operator observes a "major problem signal" requiring protracted shutdown of plant;

d_t = ***duration of spell in*** *previous* ***month***;

$d_t = 1$ "major problem spell" duration;

$d_t = 2$ refueling duration;

$d_t = 3$ operating duration.

Control variable a_t:

If $f_t < 3$, the operator has not received a major problem signal and the choice set is $A_t(x_t) = \{1, \ldots, 8\}$, given by

$a_t = 1$ close the NPP;

$a_t = 2$ refuel the NPP;

$a_t = 3$ shut down NPP (i.e., run plant at 0%);

$a_t = 4$ run the plant between 1% and 25% utilization;

$a_t = 5$ run the plant between 26% and 50% utilization;

$a_t = 6$ run the plant between 51% and 75% utilization;

$a_t = 7$ run the plant between 76% and 99% utilization;

$a_t = 8$ run the plant at 100% of its potential output.

If $f_t = 3$, the plant operator receives a major problem signal and the choice set is $A_t(x_t) = \{1, 2, 3\}$, given by

$a_t = 1$ close the NPP;

$a_t = 2$ refuel the NPP;

$a_t = 3$ run the NPP at 0% capacity.

The timing of plant signals and operating decisions is as follows: at the start of period t the plant operator knows the state r_t of the plant 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 plant operator to 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 $t + 1$. Our assumption that the NPP operator observes a signal at the start of the month summarizing the NPP's status for the rest of the month is an idealization designed so that our discrete time model could mimic the actual control process that occurs in continuous time. Given our interpretation of our DP model as an approximation of the actual 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.

Under our definition of state and control variables, a "forced outage" corresponds to the pair $(a_t = 3, f_t = 2)$, whereas a "scheduled outage" corresponds to the pair $(a_t = 3, f_t = 1)$. Not all forced outages, however, lead to complete shutdowns for the entire month. Recall that many signals represent "false alarms" that only require brief shutdowns of the plant. For example, combinations, such as $(a_t = 7, f_t = 2)$ or $(a_t = 8, f_t = 2)$, can be interpreted as the result of brief shutdowns following one or more false alarms. Forced outages can occur for a variety of reasons. Our definition of the f_t state variable does not distinguish the number or cause of the forced outage: conditional on the event that one or more forced outages occurred, we assume that all the additional information about the number and severity of these outages is captured by the unobserved state variable ϵ_t.

We distinguish, however, between outages that succeed in diagnosing and repairing problems in a short time versus protracted outages that can last for many months. We define a "major problem spell" as any continuous shutdown that lasts longer than nine months. Major problem spells are infrequent events (there are 56 such spells in our dataset) that occur for a variety of reasons, including overhaul or replacement of major reactor components (such as steam generators in a PWR), and outages mandated by the NRC due to safety concerns (such as the multi-year shutdowns of the Browns Ferry plants, and the closure of the Yankee Rowe plant following discovery of cracks in the reactor vessel). We assume that major problem spells are exogenous stochastic events, i.e., the operator lacks control over their incidence and duration.

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:

$$r_{t+1} = \begin{cases} 1 & \text{if } f_t = 3 \\ 2 & \text{if } a_t = 2 \text{ and } f_t < 3 \\ 3 & \text{if } a_t > 2 \text{ and } f_t < 3, \end{cases} \tag{4.9}$$

$$d_{t+1} = \begin{cases} d_t + 1 & \text{if } r_{t+1} = r_t \\ 1 & otherwise. \end{cases} \tag{4.10}$$

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 an NPP takes time, our model will simply estimate a parameter representing the net discounted costs involved in the plant closure as a one time charge. If the plant 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., $A_{480}(x) = \{1\}$.

The law of motion for the NPP status variable f_t is probabilistic, and its probability distribution is derived from 5 conditional probabilies denoted by p_{of}, p_{rf}, p_{om}, p_{mo}, and p_{ro} defined by

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

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

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

p_{mo} probability of coming up (i.e., resuming operation) from a major problem spell; and

p_{ro} probability of coming up (i.e., resuming operation) from a refueling outage.

The last probability, p_{ro}, is only relevant in the exogenous refueling model where refueling durations are assumed to be beyond the control of the operator. In the endogenous model this probability is derived as part of the solution to the DP problem. 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

$$p_i(x_t, a_t, t) = \frac{\exp\{g(x_t, a_t, t, \psi_i)\}}{1 + \exp\{g(x_t, a_t, t, \psi_i)\}}, \quad i = 1, \ldots, 5, \tag{4.11}$$

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, depending on whether the plant is currently in a major problem spell. If the plant is not in a major problem spell (i.e., if $f_t < 3$), then f_{t+1} is given by

$$f_{t+1} = \begin{cases} 1 & \text{with probability } (1 - p_{om})\,(1 - p_{of}) \\ 2 & \text{with probability } (1 - p_{om})\,p_{of} \\ 3 & \text{with probability } p_{om}. \end{cases} \tag{4.12}$$

If the plant is currently in a major problem spell (i.e., if $f_t = 3$), then f_{t+1} is given by

$$f_{t+1} = \begin{cases} 1 & \text{with probability } p_{mo}(1 - p_{of}) \\ 2 & \text{with probability } p_{mo}\,p_{of} \\ 3 & \text{with probability } (1 - p_{mo}). \end{cases} \tag{4.13}$$

In the exogenous refueling specification of the DP model 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 (4.13) but with p_{ro} replacing p_{mo}. The interpretation here is that if the plant was refueling last month (i.e., if $r_t = 2$) and if the operator receives the signal $f_t = 3$ at the beginning of the current month, he has learned that there are problems that will force the plant to remain in the refueling outage for the remainder of the month. On the other hand, the signal $f_t = 1$ corresponds to information that refueling has ended and no forced outages will occur during the current month. The signal $f_t = 2$ corresponds to information that refueling has ended and the operator will experience one or more forced outages, possibly associated with the cold start-up of the

NPP following the refueling outage. This interpretation requires the implicit assumption that there is zero probability of entering a major problem spell once refueling is in progress. In summary, in the exogenous refueling model the interpretation of the signal $f_t = 3$ depends on the state that the plant was in last month. If the plant was operating last month ($r_t = 3$), then $f_t = 3$ signals the beginning of a major problem spell. If the plant was in a refueling outage last month ($r_t = 2$), then $f_t = 3$ signals a continuation of the refueling outage for another month. If the plant was in a major problem spell last month ($r_t = 1$), then $f_t = 3$ signals the continuation of the major problem spell for another month.

The NPP's profit function $\pi(a_t, x_t, \epsilon_t)$ was specified in equation (4.1). Let $u(a)$ denote the level of electricity generated by the plant, given the utilization decision, a (i.e., the product of the plant's size and its capacity factor). 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

$$\mu(a_t, x_t, \phi) = \begin{cases} -\phi_d & \text{if } a_t = 1 \text{ (close the plant)} \\ -c_r(x_t, \phi_r) & \text{if } a_t = 2 \text{ (refuel the plant)} \\ p_t u(a_t) - c_o(x_t, \phi_o) & \text{if } a_t > 2 \text{ (operate the plant at level) } a, \end{cases} \tag{4.14}$$

where $c_r(x, \phi_r)$ denotes the expected cost of refueling in state x; $c_o(x, \phi_o)$ denotes the expected cost of operating a plant in state x; and ϕ_d denotes the present value of costs associated with closing and decommissioning the plant. The unknown parameter vector $\phi = (\phi_d, \phi_r, \phi_o)$ will be specified in more detail in the next section.

Also, it is impossible to identify the location and scale of the utility's profit function using only data on operating histories. The reason is simple. It is easy to see from equation (4.2) that we obtain the same optimal decision rule $(\alpha_0, \ldots, \alpha_T)$ from any monotonic affine transformation of the profit function, $\delta_1 \pi(a_t, x_t, \epsilon_t) + \delta_2$, where δ_1 and δ_2 are constants and $\delta_1 > 0$. Therefore, we must impose an arbitrary normalization of location and scale.

To simplify our model, we assume that the price of electricity is constant over time. We normalize the profit function by dividing π by the product of the plant's size and the electricity price p. The scale normalization is completed by assuming that the normalized error term ϵ_t has a standard Type I extreme value distribution. The location normalization can be imposed by assuming that $\mu(a, x, \phi) = 0$ for a prespecified state and decision pair (a, x). 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 an NPP is independent of its size. We plan to relax this assumption in future work.

5. Estimation Results

This section presents structural estimation results for the DP model presented in Section 4. Specifically, we estimate the parameter vector $\theta = (\beta, \phi, \psi)$ specifying the discount factor, and the unknown parameters of the profit function, and law of motion for the state variables. We then evaluate the ability of the estimated DP model to fit the data through a series of comparisons of predicted versus observed choice probabilities and Chi-square goodness of fit tests.

We obtained better results from the exogenous refueling specification of the DP model discussed in Section 4, so we focus on that version of the model below. We also had some difficulties in estimating the full version of the DP model that includes "major problem spells" described in Section 4. Therefore this section presents estimation results for a simplified version of the DP model where the state r_t takes on two possible values, $r_t = 1$ (previous month was part of a refueling spell), and $r_t = 2$ (previous month was part of an operating spell). The NPP state f_t continues to take 3 possible values, but in this version of the model, $f_t = 3$ corresponds to the signal that the current refueling spell will be continued for another month. Since the possibility of major problem spells is ruled out, the DP model assigns probability 0 to the signal $f_t = 3$ during an operating spell (i.e., when $r_t = 2$ and $a_t > 2$). We also removed the 56 major problem spells (defined as outages lasting more than 9 months) from our estimation sample.

Another issue is how to determine whether a given month was a refueling month or an operating month. Given that most refueling spells begin and end midway through a month, it is important to use a classification procedure that minimizes the resulting time aggregation errors. This is particularly critical for obtaining an accurate estimate of the discretized distribution of refueling outages. The procedure we adopted involved two stages. In the first stage we shifted the beginning or ending dates of a refueling spell backward or forward by a fraction of a month to create an artificial spell that starts or ends at the beginning or end of a month. This minimizes the number of months where there are both refueling and operating hours, thus minimizing the errors involved in classifying months as operating or refueling months. For example, suppose an NPP begins a refueling spell midnight on September 16^{th}, continues refueling throughout October, and completes the refueling at midnight on November 21^{th}. The total duration of the refueling spell is 1584 hours consisting of 360 hours in September, 744 hours in October, and 480 hours in November. Our recoding procedure moves 240 of the 360 refueling hours in September to November and treats the shifted refueling spell as beginning on the 21^{st} day of September and ending on the final day in November yielding 120 refueling hours in September and 720 hours in October and November.[30] In the second stage we used the shifted refueling spells together with a threshold of 360 hours per month to classify each month as an operating or refueling month. In our example, September would classified as an operating month (since the recoded refueling hours is only 120), while October and November are refueling months. The discretized duration of this recoded refueling spell is 2 months or

[30] We ignore the gain of an hour in the autumn from "falling back" from day-light savings time.

1440 hours compared to 3 months or 2160 hours for a 360 hour cutoff rule that did not shift the beginning or ending dates of refueling spells. We have estimated the DP model with and without shifting refueling dates in this way and find that our shifting procedure gives the most accurate results, although none of the results change dramatically when we adopt different discretization methods for coding the data. (We provide some evidence of the accuracy of our procedure in Table 5.6. Without our recoding procedure mean refueling durations would be overestimated by 33%.)

We did not use this recoding procedure to estimate the unknown parameters of the utility's profit function, however. Instead we define a refueling month as any month with positive refueling hours. The reason is that adopting a higher threshold (such as 360 hours) would result in months before and after a refueling spell that will be classified as operating months but whose availability is less than 100% when we subtract the downtime for refueling. In our discrete time DP model this classification would be interpreted as a decision by the operator to run the plant at lower availability in the months before and after a refueling spell, even though in continuous time the NPP may have been operated at 100% capacity immediately before and after the refueling spell. We have experimented with various discretization procedures for estimating both $\lambda(x'|x, a, \psi)$ and $\mu(a, x, \phi)$ and have found that this combination of methods works well. However none of our results change dramatically when other discretization methods are used.

5.1 Estimating the Probability of a Forced Outage after Refueling: p_{rf}

The probability of experiencing one or more forced outages in the month after the end of a refueling spell is p_{rf}. We estimate this probability separately from the probability p_{of} of experiencing one or more forced outages once an operating spell is in progress. This is because there is an unusually high rate of problems associated with the cold start of an NPP in the months after a refueling spell. We provide two sets of estimation results: an "unrestricted" model that includes plant-specific variables (and other variables designed to check the validity of our specification) and the restricted version that was used in the solution and estimation of the DP model.

Table 5.1 presents the unrestricted estimation results for ψ_{rf}, the parameters of a linear-in-parameters specification of the binomial logit model for p_{rf}. There are few observed covariates that correlate well with plant-level heterogeneity in the rate of forced outages after a refueling. The most important effect is from plant capacity: medium and very large NPPs have significantly lower rates of forced outages after a refueling than small plants (whose coefficient is normalized to zero). We included the duration of the previous operating spell to test whether the refueling spell is able to fully "regenerate" the NPP. We find no evidence that longer operating spells are associated with higher rates of forced outages after a refueling.

The results from the restricted specification for p_{rf} are given in Table 5.2. Removing the covariates reduces the log likelihood from -485 to -501. Although a likelihood ratio test rejects the hypothesis that the covariates are irrelevant in predicting forced outage rates, the inclusion of these covariates does not have large a impact on the parameters associated with duration and age. The results confirm our earlier findings of a monotonic decrease in the rate of forced outages with age, probably representing a combination of general technological improvement and plant-level learning-by-doing as discussed in Section 3. The coefficient estimates on the quadratic duration terms show that longer refueling outages decrease the probability of experiencing forced outages on startup, at least for durations up to 4 months. For refueling outages that take longer than 4 months, the longer the outage the more likely the chance of experiencing a forced outage in the month after startup.

Parameter	Estimate	Standard Error	t-statistic
Constant	2.67	$.82$	3.25
Westinghouse	$-.42$	$.30$	1.38
Combustion Engineering	$-.47$	$.33$	-1.41
General Electric	$-.57$	$.32$	-1.79
$[400, 750)$ MWe	-1.38	$.57$	-2.43
$[750, 1000)$ MWe	$-.89$	$.57$	-1.55
$[1000, \infty)$ MWe	-1.39	$.59$	-2.34
Age	$-.00341$	$.005$	$-.68$
Age^2	3.0×10^{-6}	1.6×10^{-5}	$.19$
d_t (previous spell)	$-.054$	$.03$	-1.78
d_t^2 (previous spell)	1.10×10^{-3}	8.61×10^{-4}	1.28
d_t (current spell)	$-.21$	$.25$	$-.86$
d_t^2 (current spell)	$.03$	$.02$	1.19
Scram in	$-.25$	$.19$	-1.28

$L(\hat{\psi}_{rf}) = -485.04$

Nobs$= 727$

Table 5.1 Unrestricted estimates of p_{rf}

Parameter	Estimate	Standard Error	t-statistic
Constant	$.34$	$.54$	$.63$
Age	$-.002$	$.001$	-1.95
d_t (current spell)	$-.14$	$.23$	$-.60$
d_t^2 (current spell)	$.021$	$.023$	$.85$

$L(\hat{\psi}_{rf}) = -500.96$

Nobs$= 727$

Table 5.2 Restricted estimates of p_{rf}

5.2 Estimating the Probability of a Forced Outage during Operation: p_{of}

The estimates of p_{of}, the probability of experiencing a forced outage during an operating spell, are presented in Tables 5.3 and 5.4. Table 5.3 presents the results of a fixed-effects logit in which separate dummy variables are included for each of the of 114 NPPs in our post-TMI sample (the coefficient for the Arkansas One Unit was normalized to zero). We included these fixed effects to ensure that our estimates of duration effects, specifically our preliminary findings of a bathtub-shaped pattern of the probability of forced outages as a function of the duration of the operating spell, is not a result of "spurious duration dependence" due to a failure to control for unobserved heterogeneity. Table 5.3 presents the estimates of the six key duration and age coefficients from the fixed effects logit estimation (the full model has a total of 131 parameters).

Our key finding is that the quadratic duration term in p_{of} is positive and statistically significant — even after controlling for unobserved heterogeneity. The function decreases for the first 13 months in the operating cycle and begins increasing thereafter, confirming the "bathtub hypothesis". We included the duration of the previous refueling outage to test whether the "investment" in longer refueling outages has a payoff in terms of lower rates of forced outages in the subsequent operating cycle. There is no evidence of such an effect. This is part of the reason why it's so difficult to model refueling durations as endogenously chosen by the operator. Without any quantifiable benefit to undergoing longer refueling outages, the DP model attempts to "explain" these durations by making refueling appear to be a profitable activity.

There was substantial variance in the estimated fixed effects, ranging from a low of -2.91 for the Beaver Valley 2 in Pennsylvania to a high of 1.26 for the Yankee Rowe (which was closed in 1991). Indeed, 5 of the 20 plants with the highest estimated fixed effects have been closed and several others (such as Browns Ferry 2) have been subject to extended NRC-mandated shutdowns.

Parameter	Estimate	Standard Error	t-statistic
Age	-1.74×10^{-2}	1.64×10^{-3}	-10.57
Age^2	2.50×10^{-5}	5.00×10^{-6}	4.65
d_t (previous spell)	6.01×10^{-3}	9.96×10^{-3}	$.61$
d_t^2 (previous spell)	1.00×10^{-5}	1.50×10^{-4}	$.06$
d_t (current spell)	-4.81×10^{-2}	1.62×10^{-2}	2.97
d_t^2 (current spell)	1.79×10^{-3}	8.65×10^{-4}	2.07

$L(\hat{\psi}_{of}) = -6079.15$

Nobs$= 10681$

Table 5.3 Unrestricted estimates of p_{of}

Table 5.4 presents the restricted specification for p_{of} used in the estimation of the DP model. We included only the linear term in age since the positive quadratic term in the unrestricted model implies an eventual increase in p_{of} after 348 months of operation. While this eventual upturn may have *a priori* plausibility from eventual aging effects, we cannot be sure that it is just an artifact of the quadratic specification because there are no observations on NPPs older than 369 months. When we omit the quadratic term in the restricted specification the coefficient of the linear term is reduced. The estimates of the spell duration terms are similar to those in the unrestricted model except that in the restricted model the probability of forced outages starts increasing after 9 months into an operating spell.

We estimated specifications with coefficients designed to capture the impact of the prior period availability decision interacted with a dummy variable for whether or not a forced outage occurred last period. As expected, forced outages are serially correlated events: a forced outage this month increases the probability of experiencing a forced outage next month. There is an interesting U-shaped effect of availability decisions on the rates of forced outages: forced outage rates are lower if the plant is shut down this month ($a_t = 3$) or run at full capacity ($a_t = 8$), but are higher for intermediate availability rates. This pattern may reveal the value of mid-cycle preventive maintenance outages and the value of steady operation of the NPP. However while these coefficients are highly significant in the restricted specification, the magnitude of these effects nearly disappear in the fixed effects logit results. This suggests that they may be artifacts of unobserved heterogeneity, which lead us to omit them in the final restricted specification presented in Table 5.4. However we have found that when we included these terms in our structural estimation of the DP model, the estimates of the profit function coefficients were not greatly affected. Figure 5.1 plots the estimated values of p_{of} as a function of the duration of the operating spell (d). We see that p_{of} has the classic bathtub shape and this function is shifted monotonically downward at a rate that is essentially linear in the age of the NPP.

Parameter	Estimate	Standard Error	t-statistic
Constant	$-.13$	.07	-1.89
Age	$-.0037$	.00031	-11.78
d_t	$-.052$	.014	-3.51
d_t^2	.0026	.0008	3.28

$L(\hat{\psi}_{of}) = -6546.45$

Nobs$=10681$

Table 5.4 Restricted estimates of p_{of}

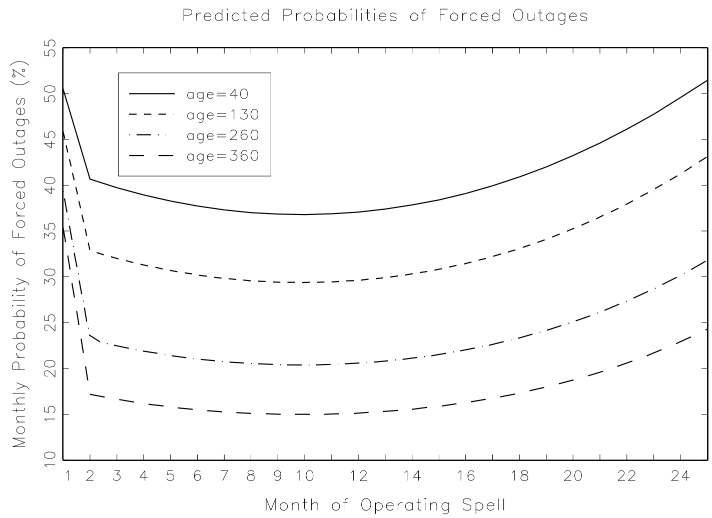

Figure 5.1 Estimated Values of p_{of}

5.3 Estimating the Probability of Ending Refueling: p_{ro}

Table 5.5 presents the estimation results for the unrestricted specification of p_{ro}, the probability of resuming operation once in a refueling spell. (This is required to estimate the exogenous refueling specification of the DP model.) The results show that General Electric's BWRs have significantly longer refueling outages than the other types of NPPs. There is some evidence that large NPPs have significantly longer refueling outages. We do not find evidence of systematic changes in the average lengths of refueling outages as the NPP ages, consistent with the results in Figure 3.12. The duration coefficients correspond to a pattern of increasing exit rates from refueling: NPPs are virtually certain to have refuelings that last between two months and 14 months. Finally, refueling outages that were entered with a scram tend to be longer than those that were entered through manual shutdowns. All of these results are reasonable.

The duration of the last operating spell has a significant impact on refueling durations: the longer the previous operating spell the longer the current refueling outage. We confirmed this effect in regressions of the duration of the refueling spell on the duration of the preceding operating spell: each additional month of operation is predicted to increase the duration of the refueling spell by 27 hours, so that an increase in the duration of the operating cycle from 12 to 24 months is predicted to increase the mean duration of a refueling outage by 1.85 months. This is in line with regression results from Stoller (1989) that predicted increasing the cycle length from 12 to 24 months for PWRs would increase refueling durations by 1.6 weeks. Although the relationship is statistically significant, it is economically insignificant, especially in relationship to other uncertain events that can have much larger impacts on the duration of refueling outages. Indeed the Stoller report concluded, “Based on surveys of utility personnel, the refueling outage duration increases from 12 to 18 month cycles and from 18 month to 24 month cycles are expected to be minimal in the future." (Stoller, 1989, p. 4-22). In view of this we decided not to attempt to explicitly model the relationship between the durations of operating cycles and the durations of subsequent refueling outages in our DP model.

Parameter	Estimate	Standard Error	t-statistic
Combustion Engineering	$-.16$	$.23$	$-.71$
Westinghouse	$.11$	$.21$	$.53$
General Electric	$-.48$	$.21$	-2.26
$[400, 750)$ MWe	$-.10$	$.36$	$-.28$
$[750, 1000)$ MWe	$-.58$	$.35$	-1.65
$[1000, \infty)$ MWe	$-.68$	$.37$	-1.86
Age	$.001$	$.003$	0.41
Age^2	$-.000007$	$.000010$	$-.68$
d_t (previous spell)	$-.04$	$.01$	-3.67
$d_t = 1$	-1.26	$.38$	-3.27
$d_t = 2$	$.99$	$.38$	2.64
$d_t = 3$	1.33	$.39$	3.43
$d_t = 4$	1.32	$.41$	3.22
$d_t = 5$	1.14	$.44$	2.60
$d_t = 6$	1.19	$.47$	2.52
$d_t \geq 7$	$-.005$	$.82$	$-.006$
$(d_t - 6)(d_t \geq 7)$	$.87$	$.39$	2.24
Scram in	$-.17$	$.12$	-1.41

$L(\hat{\psi}_{ro}) = -1183.08$

Nobs$= 2202$

Table 5.5 Unrestricted estimates of p_{ro}

Table 5.6 presents the estimation results for the restricted version of p_{ro} used to estimate the exogenous refueling specification of the DP model. Age remains insignificant even when other covariates and the Age^2 term are dropped, so we excluded it from our final model. However the general pattern of increasing exit rates from refueling is robust. There is a significant drop in the likelihood in moving from the unrestricted to the restricted model suggesting the potential importance of trying to account for the effect of heterogeneity on the duration of refueling outages. However we will ignore this evidence for now and proceed to the estimation of the homogeneous specification of the DP model. Figure 5.2 plots the estimated values of p_{ro}. The exit rates in Figure 5.2 correspond to a duration of refuelings that has a mean of 2.95 months and a standard deviation of 1.57 months. These predicted values are close to the observed mean duration of refuelings of 2.94 months and standard deviation of 1.40 months. The implied distribution of refueling durations matches the observed distribution reasonably well, as shown in Table 5.7.[31]

Parameter	Estimate	Standard Error	t-statistic
$d_t = 1$	-2.37	$.13$	-18.12
$d_t = 2$	$-.22$	$.08$	-2.82
$d_t = 3$	$.02$	$.10$	$.21$
$d_t = 4$	$-.05$	$.15$	$-.36$
$d_t = 5$	$-.23$	$.20$	-1.11
$d_t = 6$	$-.15$	$.27$	$-.54$
$d_t \geq 7$	-1.30	$.68$	-1.76
$(d_t - 6)(d_t \geq 7)$	$.81$	$.38$	2.15

$L(\hat{\psi}_{ro}) = -1219.33$

Nobs= 2202

Table 5.6 Restricted estimates of p_{ro}

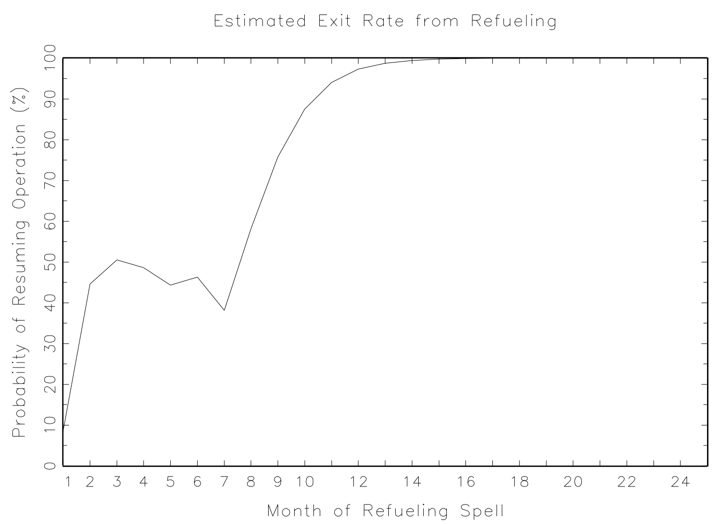

Figure 5.2 Estimated Values of p_{ro}

[31] By "observed distribution" of refueling outages we mean tabulations based on the continuous distribution computed from the Graybook's data on the number of hours spent in refueling. We used a simple continuity correction to map the continuous distribution into the discrete distribution listed as the "observed distribution" presented in Table 5.7. The histogram has bins of width 720 hours (the number of hours in a typical month) starting at 360 hours. Thus, if the number of hours in refueling were between 360 hours and 1080 hours (i.e., the bin centered on 720 hours) the refueling was classified as taking one month, and so on.

	Month of Refueling Outage											
Distribution	1	2	3	4	5	6	7	8	9	10	11	12
Actual	7.70	41.40	25.45	13.07	5.78	3.58	1.38	1.38	0.28	0.00	0.00	0.00
Estimated	8.56	40.78	25.60	12.20	5.70	3.32	1.47	1.38	0.76	0.21	0.03	0.00

Table 5.7 Comparison of Actual versus Estimated Distributions of the Duration of Refueling Outages

5.4 Estimation of the profit function, $\mu(a, x, \phi)$

We have estimated an unrestricted version of the profit function $\mu(a, d, f, \phi)$ defined in terms of a 12×1 vector of coefficients given in Table 5.8. The estimation of μ requires identifying normalizations of location and scale. The location normalization is that profit associated with shutting down the NPP immediately after a refueling outage for maintenance is zero, i.e., $\mu(3, 1, 1) = 0$. The scale normalization is accomplished by dividing the profit function π by the product of the price of electricity (assumed constant over time and the same for all plants) and the plant's size. Finally, the distribution of the unobserved components of the profit function, ϵ, has a standardized Type I multivariate extreme value distribution.

Parameter	Description
ϕ_c	expected present discounted value of costs of decommissioning NPP
ϕ_r	expected cost of refueling plant, $\mu(2, d, f) = \phi_r$
ϕ_{rf}	additional cost of refueling plant if a forced outage signal is received, $\mu(2, d, 2) = \phi_r + \phi_{rf}$
$\phi_{d,u>0}$	effect of operating cycle duration on expected profits (given positive utilization)
$\phi_{d^2,u>0}$	effect of square of operating cycle duration on expected profits (given positive utilization)
$\phi_{u=13}$	expected profit of utilization between 0 and 25%, $\mu(4, d, f) = \phi_{u=13} + d\phi_{d,u>0} + d^2\phi_{d^2,u>0}$
$\phi_{u=38}$	expected profit of utilization between 26 and 50%, $\mu(5, d, f) = \phi_{u=38} + d\phi_{d,u>0} + d^2\phi_{d^2,u>0}$
$\phi_{u=63}$	expected profit of utilization between 51 and 75%, $\mu(6, d, f) = \phi_{u=63} + d\phi_{d,u>0} + d^2\phi_{d^2,u>0}$
$\phi_{u=88}$	expected profit of utilization between 76 and 99%, $\mu(7, d, f) = \phi_{u=88} + d\phi_{d,u>0} + d^2\phi_{d^2,u>0}$
$\phi_{u=100}$	expected profit of 100% utilization, $\mu(8, d, f) = \phi_{u=100} + d\phi_{d,u>0} + d^2\phi_{d^2,u>0}$
$\phi_{u=0,f=2}$	expected profit of shutting down NPP given a forced outage signal, $\mu(3, d, 2) = \phi_{u=0,f=2}$
$\phi_{u=100,f=2}$	reduction in expected profit due to decision to operate NPP at 100% availability given a forced outage signal, $\mu(8, d, 2) = \phi_{u=100,f=2} + \phi_{u=100} + d\phi_{d,u>0} + d^2\phi_{d^2,u>0}$

Table 5.8 Definitions of Profit Function Coefficients, ϕ

Table 5.9 presents the parameter estimates of the unrestricted profit function. These estimates were computed from the two-stage partial likelihood estimator described in Section 4. The estimated standard errors have not been corrected to account for the effect of estimation error in the first stage estimates of the parameters ψ entering the transition density $\lambda(s'|s,a,\psi)$. The estimates assume a monthly discount factor of $\beta = .9999$ that corresponds to an annual real discount rate of 1.2%.[32] Overall, the DP model tells us that utilities are not myopic and that they appear to make their decisions with a high weight on the future consequences of current actions.

Parameter	Estimate	Standard Error	t-statistic
ϕ_c	24130	487	49.58
ϕ_r	$-.024$	$.19$	$-.12$
ϕ_{rf}	-3.09	$.12$	-24.96
$\phi_{u=0,f=2}$	-3.90	$.19$	-19.69
$\phi_{u=13}$	-1.51	$.13$	-11.22
$\phi_{u=38}$	$-.99$	$.13$	-7.81
$\phi_{u=63}$	$-.18$	$.12$	-1.54
$\phi_{u=88}$	1.14	$.12$	9.67
$\phi_{u=100}$	3.35	$.11$	29.32
$\phi_{u=0,f=2}$	-3.90	$.20$	-19.69
$\phi_{u=100,f=2}$	-5.31	$.11$	-50.30
$\phi_{d,u>0}$	$-.052$	$.017$	-2.96
$\phi_{d^2,u>0}$	$.0003$	$.0006$	0.52

$L(\hat{\phi},\hat{\psi}) = -10749.75$

Nobs= 13508

Table 5.9 Unrestricted structural estimates of μ

The parameter ϕ_c, representing the cost of closing and decommissioning, is positive. Is this reasonable? The value of this parameter must be interpreted relative to our normalization that the single period profit corresponding to shutting down the NPP is 0. Since the profit associated with a shutdown will be negative, the estimate of ϕ_c implies that utilities perceive a high cost of decommissioning, as we will show subsequently. In this paper we do not attempt to relate the estimates of the profit function and implied costs of decommissioning to real data. In this paper we will evaluate the DP model in terms of the credibility of its predictions and its ability to fit the data. We leave the analysis of the accuracy of its estimates of decommissioning costs and operating profits to future research.

The cost of refueling, ϕ_r, is estimated to be a small and statistically insignificant (subject to our caveat about the estimated standard errors). The results show that utilities do not perceive a refueling outage to be significantly more costly than an maintenance outage (which is 0 by our identifying normalization). This suggests that the main

[32] We experimented with different discount factors and found the likelihood was basically flat (with a small positive slope) for $\beta > .999$. Thus we are unable to identify β precisely although since the likelihood function falls rapidly for $\beta < .99$ we can easily reject the hypothesis that utilities have high discount rates.

cost of a refueling outage or maintenance outage is the opportunity cost of lost power generation rather than the costs of the maintenance activities that is being performed during these outages. This is something that we hope to verify in a future analysis when we include O&M costs.[33]

The estimates of the parameters $\phi_{u=13}$, $\phi_{u=38}$, and $\phi_{u=63}$ are all negative. This could be a counter-intuitive result: utilities perceive a higher profit from no availability than from some availability. (See Figure 3.1 on the frequency of "some"availability.) These coefficients could represent utilities' expected negative profits from unplanned shutdowns during the month. Given that startups and shutdowns create stress on the reactor, it is reasonable to conclude that utilities perceive that the costs of unexpected outages to be greater than the cost of an expected maintenance outage.

We acknowledge that these negative coefficients could also reflect possible misspecification of the DP model. For example, we assumed that utilities are risk neutral. If utilities are risk averse, the negative coefficients on these coefficients could signal their aversion to erratic operation of their NPPs, given that most NPPs are designed to satisfy base load. Another possible source of misspecification is that our DP model may not adequately capture the future benefits of preventive maintenance outages, which may lead to lower rates of forced outages or less serious outages in future months. Our model does capture the "regeneration" effect of a refueling outage via the bathtub shaped probability of forced outages described earlier. However we did not attempt to model a reduction in the probability of forced outages after mid-cycle preventive maintenance outages. The resolution of whether the negative coefficients on $\phi_{u=13}$, $\phi_{u=38}$ and $\phi_{u=63}$ reflect specification error or an aversion to stopping and starting the NPP is a topic we leave to future research.

The estimated coefficients of $\phi_{u=88}$ and $\phi_{u=100}$ are positive and statistically significant. The positive estimate of $\phi_{u=88}$ seems to show that at a sufficiently high level of availability the value of power generation outweighs the costs of a partial outage. Of course, the highest profit is associated with continuous 100% availability of the NPP and this is clearly reflected in our estimates.

The estimated coefficients of the effect of duration on the profits associated with positive availability of the NPP, $\phi_{d,u>0}$ and $\phi_{d^2,u>0}$, show that expected costs of availability increase linearly with the duration of the operating spell (the quadratic term is insignificantly different from 0). This finding may show that in addition to the effect of within-cycle deterioration on the *frequency* of unplanned events (as reflected, for example, in the increasing probability of forced outages after 12 months of operation), the *severity* of unplanned events arising from this deterioration also tends to increase with the duration of the operating spell. For example, the fraction of outages from "false alarms" may tend to decrease and the fraction of outages associated with more costly problems, such as fuel rod failures, may

[33] O&M costs are reported on an annual basis and can be found in FERC Form 1; see EIA (1988). So monthly data is not readily available. Also, nuclear fuel costs are accounted as a capital expense because of the length of time the utility holds the uranium in inventory (unless the utility leases it from nuclear fuel services company). Therefore, fuel costs are not associated with the refueling outage as such.

tend to increase with the operating spell. If this is the case, it is consistent with our finding of an approximately linear rate of increase in operating costs with the duration of an operating spell.

The final three coefficients in our specification, ϕ_{rf}, $\phi_{u=0,f=2}$ and $\phi_{u=100,f=2}$ were included to directly capture the "severity effects" of forced outages. We interpret these coefficients as reflecting the utility's expectation of the reduction in current operating profits associated with three courses of action — refueling, plant shutdown, and continued operation at 100% availability — conditional on receiving one or more forced outage signals. All coefficient estimates have large negative values, suggesting that utilities impute high costs to forced outages. The fact that $\phi_{u=100,f=2} = -5.3$ means that the potential costs of continuing to run at 100% availability after receiving a forced outage signal far outweigh the benefits of the electricity revenue generated. Similarly, the negative values on ϕ_{rf} and $\phi_{u=0,f=2}$ could imply that utilities perceive the costs of planned outage are much higher than forced outages. These negative values could reflect that fact that maintenance and repair activities conducted to repair damage arising from forced outages is more expensive. Finally, the large negative values of these coefficients may also reflect potential specification error to the extent that a forced outage is indicative of long lasting problems that can't be completely fixed in the same month as they occur.

5.5 Evaluation of DP Model's Predictions of Optimal NPP Operations

The remainder of this section will focus on judging the DP model in terms of the reasonableness of its predictions about NPP operations. We will show that our parsimonious DP model accurately predicts the operations of U.S. NPPs, suggesting that utilities may indeed be using optimal strategies for running their plants. Following Rust (1995), the hypothesis that utilities are behaving optimally is untestable in the absence of strong restrictions on their beliefs λ and objective function μ. Rust's results show that essentially any behavior pattern can be "rationalized" as optimal for some choice of (β, μ, λ). However there is no guarantee that any choice of (β, μ, λ) that does rationalize behavior will be "reasonable." To the extent that our estimates of (β, μ, λ) are regarded as reasonable, our results support the hypothesis that utilities are behaving optimally.

We begin with Table 5.9 that compares the observed (or non-parametric) estimates of the conditional choice probability function $P_t(a|x)$ (labelled "NP") and the predicted values $P_t(a|x, \hat{\phi})$ from the estimated DP model (labelled "DP"). The NP estimate can be computed as a simple histogram using the observations in conditioning cell (t, x). The DP estimate is computed using the logit formula for $P_t(a|x, \hat{\phi})$ given in equation (4.4) of Section 4 and the estimated coefficients $\hat{\theta}$ given in Table 5.8.

Unfortunately, there are many more conditioning cells (t, x) than we have data: there are a total of $8,873$ distinct (t, x) combinations in our dataset whereas in the DP model there are $72,000$ such cells ($72,000 = 480 \times 25 \times 3 \times 2$, where there are $T = 480$ possible ages, 25 possible values for the duration variable d_t, 3 possible values for the NPP signal variable f_t, and 2 possible values for the NPP state variable r_t). With a total of 13,508 reactor-months we have an average of only 1.5 observations per cell that is too small to estimate the 7 unknown probabilities $P_t(a|x)$ in each (x,t) cell. We must aggregate these elemental cells to make statistically reliable comparisons. Given a subset X of (x,t) cells, it is easy to compute the parametric and non-parametric estimates of $P(a|X)$ by sample enumeration:

$$\begin{aligned} \hat{P}(a|X) &= \int_{x\in X} \hat{P}(a|x)\hat{F}(dx|X) \quad \text{(NP)} \\ P(a|x,\hat{\theta}) &= \int_{x\in X} P(a|x,\hat{\theta})\hat{F}(dx|X) \quad \text{(DP)} \end{aligned} \tag{5.1}$$

where $\hat{F}(dx|X)$ is the non-parametric estimate of the conditional probability distribution of x given X. It is equal to $N(dx)/N(X)$, where $N(x)$ is the number of observations in cell dx and $N(X)$ is the total number of observations in cell X.[34]

Table 5.10 presents the comparisons of the NP and DP estimates for the full sample, i.e., when there is a single group X containing all possible (t,x) combinations in the data. The results show that the DP model fits the observed choices closely. The DP model accurately predicts the rare event of NPP closure as well as the very frequent event of 100% availability.

Parameter	NP	DP
$a = 1$ (close)	.037	.037
$a = 2$ (refuel)	20.49	20.47
$a = 3$ ($u = 0\%$)	3.29	3.25
$a = 4$ ($u = 13\%$)	1.52	1.54
$a = 5$ ($u = 38\%$)	2.54	2.58
$a = 6$ ($u = 63\%$)	5.69	5.78
$a = 7$ ($u = 88\%$)	21.36	21.68
$a = 8$ ($u = 100\%$)	45.08	44.66
Number of Cells	8873	
Number of Observations	13508	

Table 5.10 Predicted versus Actual Choice Probabilities: Full Sample

[34] Given a (possibly random) partition of the x cells, one can compute an asymptotic Chi-squared goodness of fit test that accounts for the covariates as described by Andrews (1988). However we do not present these statistics because they require corrected standard errors for the $\hat{\phi}$ coefficients. However we feel that direct comparisons the NP and DP estimates is more useful since it shows exactly where the DP model does and doesn't fit well and illustrates the predictions of the model — something a single Chi-squared statistic is incapable of doing.

Table 5.11 presents comparisons of the NP and DP estimates for the subsample of NPPs that are in state $r_t = 1$, i.e. they were in a refueling spell in the preceding month. In the exogenous specification of the DP model, if the operator receives a signal $f_t = 3$, the refueling outage will continue for the rest of the current month (whereas the signal $f_t = 1$ or $f_t = 2$ corresponds to information that the refueling outage is over and the NPP is able to resume operations in the current month). In the case where $f_t = 3$ the operator's choice set is $A_t(x_t) = \{1, 2\}$ corresponding to the option of closing the NPP or continuing the refueling whereas if $f_t = 1$ or $f_t = 2$ the operator's choice set is $A_t(x_t) = \{1, 3, 4, 5, 6, 7\}$, i.e., the full choice set except for the refueling option. The results show that the DP model does a good job in predicting NPP availability once it comes up from a refueling. Together with our accurate estimates of p_{ro}, the probability of remaining in a refueling outage (see Table 5.7), the DP model is able to accurately capture both the duration of refueling outages and the potential problems operators experience in the first month after a cold start of the NPP. This is verified in separate tabulations (not shown): the DP model closely predicts the operator's availability choices for the subsample $(r_t = 1, f_t = 1)$ (i.e., for NPPs that don't experience any forced outages in the month after startup) and also for the subsample $(r_t = 1, f_t = 2)$ (i.e., for NPPs that do experience one or more forced outages in the month after startup).

Parameter	NP	DP
$a = 1$ (close)	0.00	$.032$
$a = 2$ (refuel)	74.19	74.16
$a = 3$ ($u = 0\%$)	0.99	0.54
$a = 4$ ($u = 13\%$)	0.64	0.66
$a = 5$ ($u = 38\%$)	1.13	1.10
$a = 6$ ($u = 63\%$)	2.24	2.47
$a = 7$ ($u = 88\%$)	9.55	9.28
$a = 8$ ($u = 100\%$)	11.25	11.75
Number of Cells	1611	
Number of Observations	2817	

Table 5.11 Predicted versus Actual Choice Probabilities: Refueling Spells

Table 5.12 presents comparisons of the NP and DP estimates for the subsample of NPPs that are in state $r_t = 2$, i.e., plants that were in an operating spell in the previous month. When an NPP is in an operating spell the operator has the option of shutting down the NPP for a mid-cycle outage (i.e., $a_t = 3$) or of shutting down the NPP for a refueling outage (i.e., $a_t = 2$). We can see from Table 5.12 that the DP accurately predicts both probabilities. Of course, these are averages of the conditional probabilities that vary with the duration of the operating spell, d_t. (We show below that the DP model also succeeds in predicting how these probabilities change as a function of d_t.)

Parameter	NP	DP
$a = 1$ (close)	.047	.037
$a = 2$ (refuel)	6.34	6.32
$a = 3$ ($u = 0\%$)	3.89	3.97
$a = 4$ ($u = 13\%$)	1.75	1.77
$a = 5$ ($u = 38\%$)	2.91	2.97
$a = 6$ ($u = 63\%$)	6.60	6.65
$a = 7$ ($u = 88\%$)	24.47	24.95
$a = 8$ ($u = 100\%$)	53.99	53.33
Number of Cells	7262	
Number of Observations	10691	

Table 5.12 Predicted versus Actual Choice Probabilities: Operating Spells

Table 5.13 illustrates an aspect of NPP operations that the DP model does not do a good job of capturing: "fine tuning" shutdowns in the second month after a refueling outage. The table compares the NP and DP estimates for the subsample of NPPs with $(r_t = 2, d_t = 1)$, i.e., plants that have completed the first month of an operating spell. The table shows that the DP model substantially underpredicts the probability of a non-refueling shutdown (i.e., decision $a_t = 3$) and overpredicts the probability of running the NPP at 100% availability (i.e., decision $a_t = 8$). It is common practice in the industry to bring up an NPP, run it for a test period to "shake down" any problems in the NPP that may have arisen from repairs that occurred in the last refueling outage, and then shut down the NPP to correct any problems. The present version of the DP model does not have sufficiently detailed state variables to capture operators' motivations for undertaking these fine tuning outages, although it is possible that by using more detailed data on reasons for outages from the Graybook we could formulate and estimate a more detailed DP model that could capture this effect in the future.[35]

We now turn to a demonstration of the ability of the DP model to track changes in NPP operations as a function of duration of the operating cycle and the age of the reactor. Figure 5.3 plots the predicted versus actual availability factors for NPPs as a function of the duration since last refueling. The top panel presents a comparison of what might be called unconditional availability factors since it includes the loss in availability due to refueling outages and plant closures as well as mid-cycle shutdowns. The figure shows that the DP model succeeds in tracking the decrease in availability factors as the time since last refueling increases, and it also shows that the model tracks the large effect of forced outages on NPP availability as noted above. The second panel plots what we might call the unconditional availability factor or "availability factor between refuelings" (AFBR) that conditions on the event the NPP was not

[35] An alternative procedure could be to introduce an additional state variable, a special "startup problem indicator" that could be treated as an additional unobserved state variable. A simpler, but admittedly *ad hoc* procedure would be to introduce a dummy variable for the combination $(a_t = 8, d_t = 1)$ that could be treated as the increase in expected operating costs from the discovery of special problems after startup of the NPP.

Parameter	NP	DP
$a = 1$ (close)	0.00	2.4×10^{-6}
$a = 2$ (refuel)	0.22	0.04
$a = 3$ ($u = 0\%$)	14.46	3.26
$a = 4$ ($u = 13\%$)	2.89	1.97
$a = 5$ ($u = 38\%$)	4.12	3.30
$a = 6$ ($u = 63\%$)	6.45	7.40
$a = 7$ ($u = 88\%$)	28.36	27.78
$a = 8$ ($u = 100\%$)	43.49	56.23
Number of Cells	627	
Number of Observations	899	

Table 5.13 Predicted versus Actual Choice Probabilities: First Month of Operating Spell

shut down for refueling or a permanent plant closure. As we can see, AFBR shows very little tendency to decline with duration since last refueling. The DP model does a good job of tracking the actual AFBR's for the first 16 months of the operating cycle, but does not track the wild gyrations in the AFBR's over the 18 to 25 month interval. As noted previously there are very few observations in this latter range, so we have little reason to believe that these wild gyrations represent anything more than estimation error from extremely small samples. The results presented here do seem to conflict with the graphs of the very rapid decline in NPP availability factors given in figure 3.9 of section 3. One reason the NP estimates of the AFBR given in figure 5.3 differs from that given in figure 3.9 is that the latter is based on the exact availability factors whereas the later is given by the product of the NP estimates of the choice probability and the vector of midpoints of our discrete utilization intervals, i.e. $(.13, , .38, .63, .88)$. When there are many observations, the fact that AFBR's are approximately uniformly distributed within each of the subintervals means that taking the midpoints is a good approximation (see figure 3.1), so figure 5.3 and 3.9 in fact match very closely over the duration range $[1, 18]$. However over the range $[19, 25]$ where there are very few observations, the "slosh" of observations within these intervals can be significant, so the two estimates can differ significantly.

We conclude that the rapid declines in AFBR's beyond 18 months of operation in figure 3.9 are artifacts of very small samples. The structural estimation results indicate that it is possible for NPPs to have very long operating spells without a rapid escalation in time lost due to planned and unplanned outages, a result consistent with the overall conclusion of the Stoller report, which investigated the relationship between CFBR (capacity factor between refuelings) and operating cycle duration by regression techniques. They concluded that "Capacity factors during the operating portion of the cycle (i.e. between refuelings) have improved with longer fuel cycles." (Stoller, 1989, p. 3-7). Their regression analysis predicts that for CFBR would increase from 81.3% to 83.5% in moving from 12 month cycles (250 effective full power days) to 24 month cycles (500 EFPD) for PWRs, and from 73.9% to 78.9% for BWRs. Our results

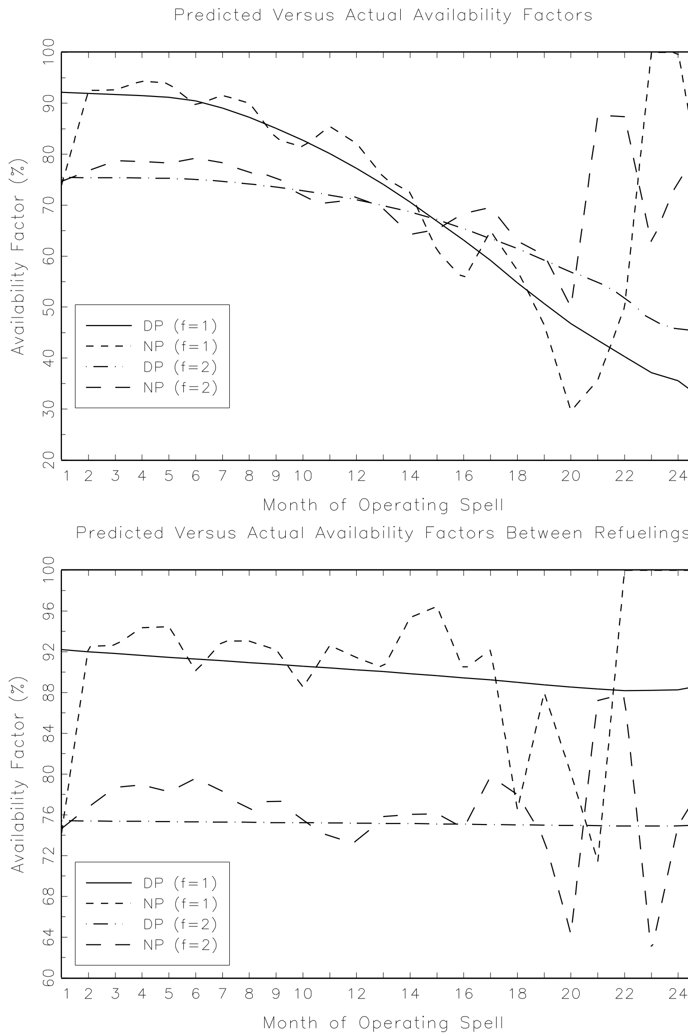

Figure 5.3 Predicted versus Actual Availability Factors

predict a slight decline in AFBRs, but this could be consistent with increasing CFBRs if NPPs tend to operate their reactors below MDC early in the operating cycle and gradually increase power to MDC later in the cycle.

Figure 5.4 demonstrates how the DP model is able to track the changes in the probabilities of refueling, non-refueling shutdowns, and the fraction of time spent in planned and forced outages as a function of the duration of the operating spell. As expected the exit rate into a refueling outage increases monotonically with duration. The DP model correctly predicts that the probability of initiating a refueling outage in a month where there are one or more forced outages is uniformly lower than in months where there are no forced outages. The second panel compares predicted versus actual probabilities of action $a_t = 3$, i.e. the decision $u = 0$. Here we see that the DP model badly underpredicts the probability of a plant shutdown after the first month into the operating cycle, a finding we previously pointed out in

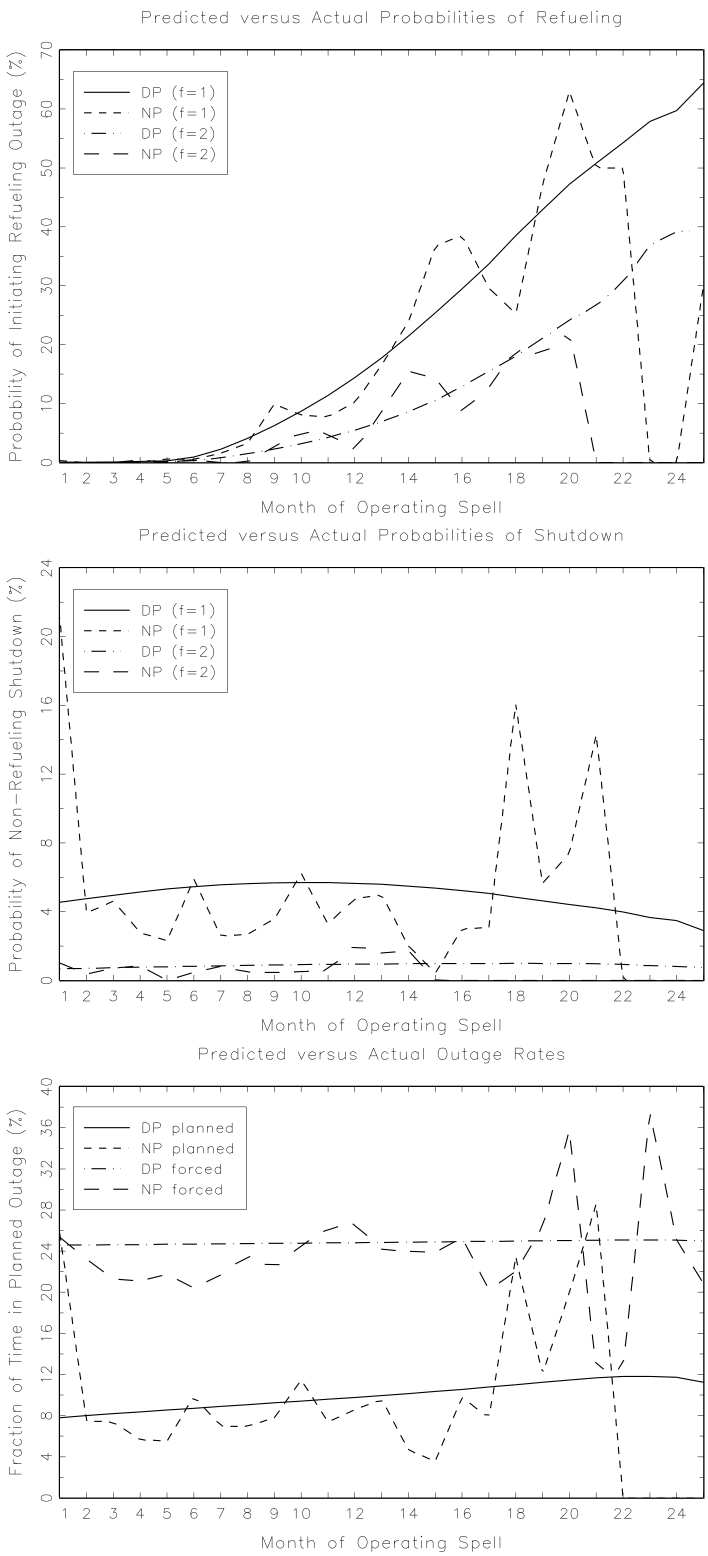

Figure 5.4 Predicted versus Actual Probabilities of Refueling, Non-Refueling Shutdowns, and Planned and Forced Outages

our discussion of "fine tuning shutdowns" in table 5.13 above. Interestingly the probability of a fine tuning shutdown is much higher when there are *no* forced outages than when there are forced outages in the first month following the refueling spell, a result that seems counterintuitive. We also see that the probability of $u = 0$ is uniformly lower when $f = 2$ than when $f = 2$, a result that also seems counterintuitive. We can see that the DP model tries to strike a balance in its prediction errors, underpredicting the probability of $u = 0$ when $d = 1$ but slightly overpredicting the probability of $u = 0$ when $d > 1$. When $f = 2$ the DP model tracks the probability of $u = 0$ very closely for all values of d. The final panel shows that the DP model is able to accurately track the loss in plant availability due to planned and forced outages. The values plotted in these graphs are conditional on the event that a forced outage did or did not occur. Thus to get the overall rates that are comparable to the unconditional values plotted in figure 3.9 of section 3 we need to multiply the values given in figure 5.4 by the probability of one or more forced outages occurring (in the case of time lost due to forced outages), or its complement (in the case of time lost due to planned outages). When we do this (not shown) we find that the DP model continues to track actual outage rates very closely, a result that follows from the accuracy of our bathtub shaped estimates of the probability of forced outages, p_{of} given in table 5.4 above.

We now turn to a discussion of the ability of the DP model to track changes in NPP operations as a function of age. Figure 5.5 presents a comparison of predicted versus actual AFBR's as a function of plant age. Consistent with our earlier findings in section 3, there is virtually no change in these rates as the plant ages. The only major change occurs after the first year of operation: availability factors tend to be lower for special reasons associated with the initial startup of the plant. The DP model is unable to account for this effect, and due to lack of data on the first year of operation, the plots in figure 5.5 have simply set AFBR to 0 in the first year.

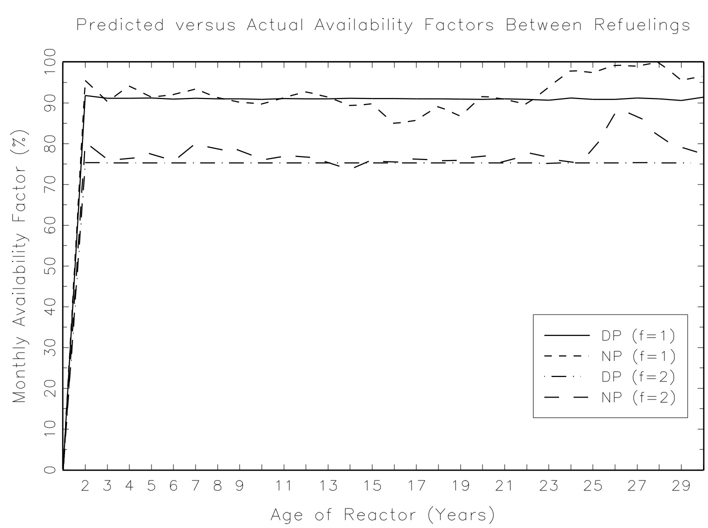

Figure 5.5 Predicted versus Actual Availability Factors Between Refuelings as a Function of Plant Age

Figure 5.6 compares predicted versus actual monthly refueling probabilities (averaged over all months in the operating cycle) as a function of the age of the reactor. We see that there is considerable year to year variation in these rates, an effect that is not due so much to behavioral variations but rather large year to year sampling variations in the conditioning variables (r_t, f_t, d_t). The main message to take away from figure 5.6 is that the DP model correctly predicts that there is no systematic trend in the probability of refueling (or in the fraction of time spent in refueling) as a function of reactor age.

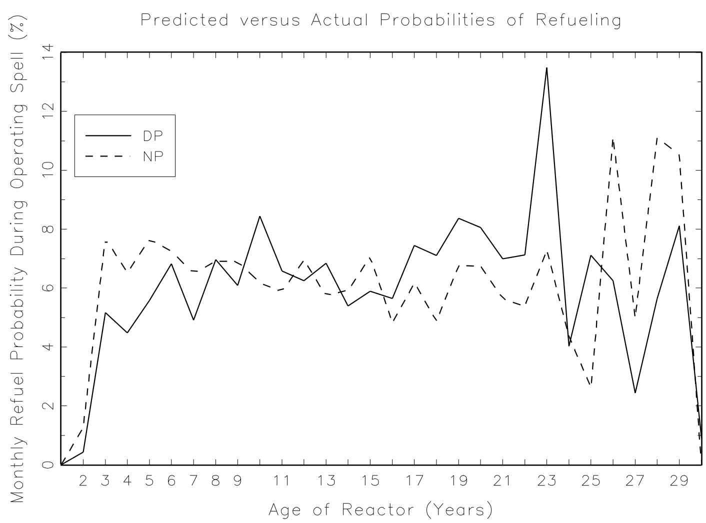

Figure 5.6 Predicted versus Actual Probabilities of Refuelings as a Function of Plant Age

We conclude our analysis of the predictions of the DP model by plotting the estimated value functions and the DP model's predictions of the probability of plant closure as a function of plant age. Figure 5.7 plots the (normalized) value functions as a function of duration of the operating spell and the age of the reactor. The first panel of figure 5.7 plots the choice-specific value functions $v_t(a_t, x_t)$ for a NPP with $t = 130, r_t = 2$, and $f_t = 1$ as a function of duration, d_t. The upper envelope of these choice-specific value functions defines the operator's optimal operating strategy as a function of duration of the operating spell (ignoring the effect of the unobserved state variables ϵ_t). We see that the model predicts that the optimal strategy is to operate the plant at 100% for the first 19 months of the operating spell and refuel thereafter. Thus, if we ignore the transitory effects of the unobserved state variables, our DP model actually predicts that NPP's have 19 month "planned" operating cycles. It is easy to verify that the presence of unobserved state variables will typically cause plants to begin refueling outages earlier than planned, which corresponds to actual operating spells being several months shorter than the planned operating cycle, a result consistent with the differences between actual and planned refueling durations noted in section 3. The value function for the decision $a = 7$ (i.e. run the plant at 88% availability) is only slightly higher than the value function for the decision $a = 3$ (i.e. shut down the plant for preventive maintenance). We interpret this finding as reflecting the strong aversion NPP operators have for start/stop operation of their plants. Note that for the first 9 months of the operating cycle the value of a non-refueling shutdown is substantially higher than the value of refueling. However after 9 months the value of refueling is basically flat as a function of d whereas the value of a non-refueling shutdown is monotonically declining in d. This explains why the probability of $u = 0$ declines after 9 to 10 months into the operating spell.

The second panel of figure 5.7 plots the value functions for the same configuration of state variables as in the top panel except that now $f_t = 2$. Now we see that the decision $a = 7$ (i.e. run plant at availability factor 88%) is the optimal decision for $d \in [1, 24]$ and $a = 2$ (i.e. refuel plant) is optimal for $d \geq 25$. Notice that the value function corresponding to $u = 100\%$ lies uniformly below the value function corresponding to $u = 13\%$ but uniformly above the value function corresponding to $u = 0\%$. We can interpret this as reflecting the potential damage that could be caused by insisting on running a reactor at 100% capacity after receiving a forced outage signal.

The final panel of figure 5.7 plots the value function as a function of age, setting the variables $d = 5$, $f = 1$ and $r = 2$ (i.e. a NPP in the fifth month of an operating spell that has not received any forced outage signals). The value functions all have a quadratic shape with a maximum at about 20 years, or halfway through the reactor's licensed lifespan. The reason why the value function decreases over the last half of the NPP's lifespan is easy to explain: it is simply the "horizon effect" that the level of expected discounted profits from running the NPP decreases as the number of remaining years in the plant's operating license decreases. But how can we explain the finding that the value function *increases* over the first half of the plant's lifespan? This turns out to be a result of the steady decrease in the rate of forced outages, an effect we have attributed to a combination of "learning by doing" and general technological progress such as improvements in fuel reliability discussed in sections 2 and 3). The flat line at the bottom of the graph is the value associated with decommissioning, assumed to be independent of the age of the plant. The quadratic shape of the value function implies, therefore, that the risk of permanent plant closure is highest at the beginning and end of the plant's operating license and lowest during the middle of its operating license.

We conclude with figure 5.8 which plots the predicted probabilities of plant decomissioning in the last few years of the plant's operating license. We see that the annualized probability of plant closure is virtually 0 until just a few years prior to the expiration of the 40 year operating license. The probability of closure is highest for plants that have experienced forced outages, and increases with the duration since last refueling. Overall, we conclude from figure 5.8 that our initial version of the DP model provides little cause for concern about a rash of premature NPP closures. However we caution the reader that this prediction could change dramatically in a version of the DP model that accounts for major problem spells: we would expect a significantly higher probability of plant closure upon entry into such a spell. Although major problem spells are relatively rare events, the fact that they occur much more frequently than plant closures suggests that the actual risk of plant closure may be significantly higher than our initial simplified version of the DP model would predict.

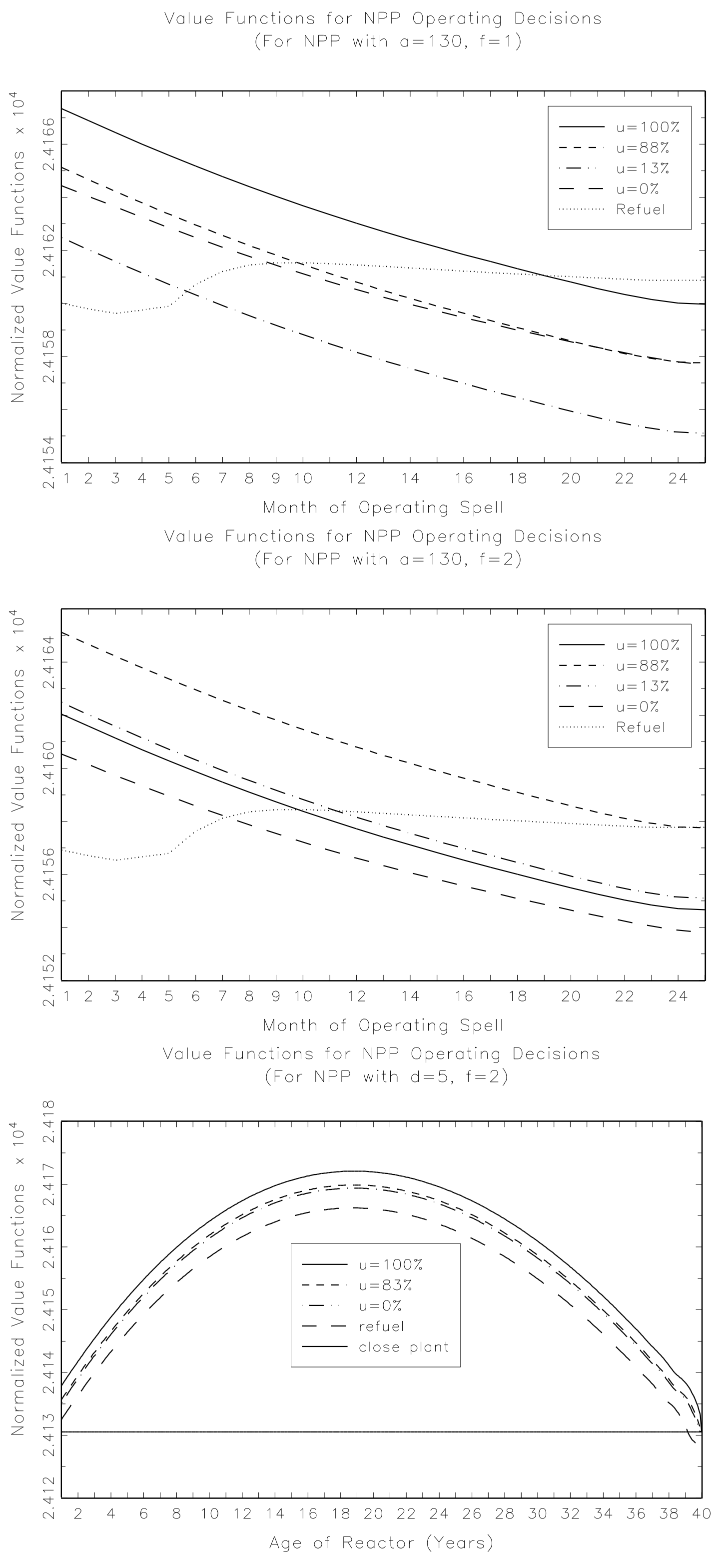

Figure 5.7 Value Functions From DP Model

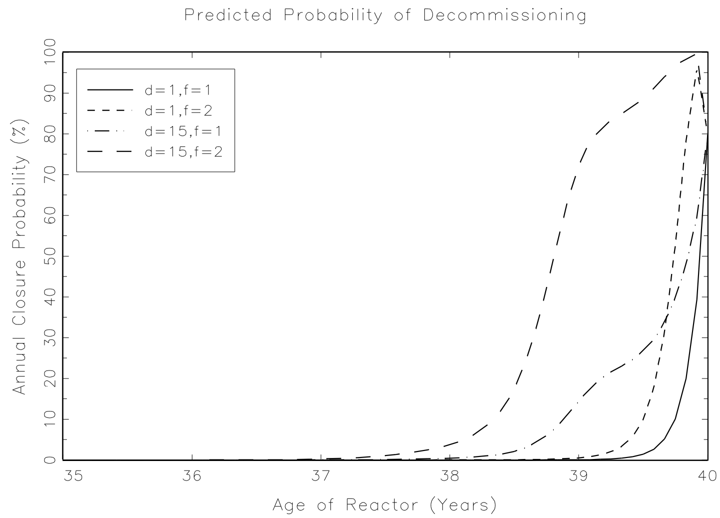

Figure 5.8 Predicted Probabilities of Plant Closure

6. Conclusions

This paper has developed a parsimonius dynamic programming model that is successful in capturing the main features of U.S. nuclear power plant operations. The results show that plant operators are extremely farsighted and highly averse to stop/start reactor operations. Utilities also impute very high costs to forced outages, which may be an indication that they are more responsive to stricter NRC safety regulations in the aftermath of the TMI accident. The DP model predicts that the optimal operating policy of utilities in this post-TMI regulatory environment is to adopt 18 to 19 month planned operating cycles, a result consistent with the observed shift in planned operating cycles from 12 months in the pre-TMI era to 18 months in the post TMI era. We are encouraged by the success of our preliminary efforts in estimating this model: in future work we expect to develop more elaborate versions of the DP model that overcome some of the limitations of the present model. Our planned extensions include:

1. estimation of a 3 state version of our model that includes “major problem spells” as described in section 4.

2. estimation of a more fine-grained daily model of NPP operations or a continuous time model along the lines of An 1992 in order to minimize problems of time aggregation and discretization error encountered in the estimation of the monthly model developed here

3. incorporation of auxiliary data on operating and maintenance expenses and decomissioning costs to assist in specification testing of the model and in the identification of technological cost factors from subjective “goodwill” and “safety” components of the cost function

4. incorporation of observed heterogeneity such as plant capacity and NSSS variables and the effects of FAC’s and variations in local regulatory policies for electric utilities

5. inclusion of the price of electricity as an additional state variable

Once these extensions are complete we plan to use the DP model to predict the impact of changes in NRC regulatory policy and to make long term predictions of the net electrical generation and retirement rates of U.S. NPPs. In particular in a future paper we will use the model to address the question raised in the introduction: was the shift in operating NPP practices in the period after TMI a result of stricter NRC regulations or other factors? We expect our model will be useful in providing insights into a number of other questions, such as the costs and benefits of alternative regulatory policies and the impacts of carbon taxes on long term operating and retirement decisions of U.S. NPPs.

7. References

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

Bamford, W.H., Cipolla, R.C., and Jaske, C.E. (eds.) (1991), *Service Experience in Operating Plants 1991* (PVP Vol. 221), New York: American Society for Mechanical Engineers.

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

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

David, P.A., and Rothwell, G.S. (1995), “Measuring Standardization: An Application to the American and French Nuclear Power Industries,"*European Journal of Political Economy* (forthcoming).

David, P.A., Rothwell, G.S., and Maude-Griffin, R. (1994), “Learning by Accident?: Reductions in the Risk of Unplanned Outages in U.S. Nuclear Power Plants After Three Mile Island,” update of Discussion Paper 248, Center for Economic Policy Research, Stanford University.

David, P.A., Rothwell, G.S., Maude-Griffin, R., and Sturm, R. (1991), “European Nuclear Power Plants and Their Less Reliable American Cousins: International Differences in the Distributions of Reactor Operating Spell Durations,"Discussion Paper 273, Center for Economic Policy Research, Stanford University.

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.

Driscoll, M.J., Downar, T.J., and Pilat, E.E. (1990), *The Linear Reactivity Model of Nuclear Fuel Management*, La Grange Park, Illinois: American Nuclear Society.

Dubin, J. and Rothwell, G.S. (1989), “Risk and Reactor Safety Systems Adoption,” *Journal of Econometrics*, 42, 201-218.

Egan, M.R. (1984), “Elements of Nuclear Reactor Fueling Theory,” *Progress in Nuclear Energy*, 14-3, 313-360.

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

Energy Information Administration (EIA) (1988), *An Analysis of Nuclear Power Plant Operating Costs*, Washington, D.C.: U.S. Department of Energy (DOE/EIA-0511).

Energy Information Administration (EIA) (1991), *Electric Plant Costs and Production Expenses 1989*, Washington, D.C.: U.S. Department of Energy (DOE/EIA-0455).

Energy Information Administration (EIA) (1992), Washington, D.C.: U.S. Department of Energy. (DOE/EIA-

Forrest, L.R. Jr., and Deutsch, T.R. (1988), “Cost Savings from Extended Life Nuclear Plants,” *Proceedings of the Topical Meeting of Nuclear Power Plant Life Extension*, Snowbird, Utah.

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.

National Research Council, (1992), *Nuclear Power: Technical and Institutional Options for the Future*, Washington, D.C.: National Academy Press.

Nero, A. V. Jr. (1979) *A Guidebook to Nuclear Reactors* University of California Press, Berkeley.

Pasqualetti, M.J., and Rothwell, G.R. (eds.) (1991), *Nuclear Decommissioning Economics,* special Issue of *The Energy Journal*.

Pershagen, B. (1989) *Light Water Reactor Safety* Pergamon Press, Oxford.

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. (1989), "Stop and Start: A Duration Analysis of Nuclear Reactor Operations," *Proceedings of the International Association of Energy Economists' Annual North American Conference*, pp. 309-317.

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

Rothwell, G.S. (1993), "Comparing Boiling and Pressurized Water Reactor Productivity in the United States: 1975-1990" in *Energy Systems and Ecology: Proceedings of an International Conference*, eds. J. Szargut, *et al*, Krakov, Poland: American Society of Mechanical Engineers.

Rothwell, G.S. (1994), "U.S. Nuclear Power Policy to the Year 2000,"in *Nuclear Power at the Crossroads*, eds. G. Hinman and T. Lowinger, International Research Center for Energy and Economic Development.

Rothwell, G.S. (1995), "Organizational Structure and Expected Output for Nuclear Power Plants, *Review of Economics and Statistics*, (forthcoming).

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.

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

Silvennoinen, P. (1976), *Reactor Core Fuel Management*, Pergamon Press, Oxford.

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).

Sturm, R. (1994), "A Structural Economic Model of Operating Cycle Management in European Nuclear Power Plants" *European Economic Review*.

U.S. Nuclear Regulatory Commission (1977), *Nuclear Power Plant Operating Experience, 1974-1975* (PB-265 794).