Journal of Risk and Uncertainty, 13:175–198 (1996)
© 1996 Kluwer Academic Publishers

Learning by Accident? Reductions in the Risk of Unplanned Outages in U.S. Nuclear Power Plants After Three Mile Island

PAUL A. DAVID
Stanford University and All Souls' College, Oxford

ROLAND MAUDE-GRIFFIN
Stanford University

GEOFFREY ROTHWELL
Stanford University

Abstract

This study uses a Cox proportional hazards model to analyze changes in the risk of unplanned outages in U.S. nuclear power plants after the Three Mile Island (TMI) accident. The unplanned outage hazard is related to safety by the fact that most such outages begin with unplanned reactor scrams. These place extreme stresses on plant equipment, increasing the risk of serious accident. The estimates indicate that the Nuclear Regulatory Commission (NRC)-led efforts to improve nuclear plant safety after TMI were followed by substantial reductions in the risk of unplanned outages.

Key words: nuclear power plant, nuclear reactor safety, proportional hazards

JEL Classification: L94, L51, Q40, C41

The worst accident in the history of U.S. commercial nuclear power occurred in the General Public Utilities' Three Mile Island (TMI) Unit 2 reactor during March of 1979. The accident partially melted the reactor core, produced external releases of radiation, and resulted in billions of dollars of plant losses, cleanup costs and litigation (Perrow, 1984; Cunningham, 1989; Lipford, Cole, and Friderichs, 1991). The TMI accident was, however, a catalyst for major efforts by the nuclear utilities and their regulators to improve nuclear plant safety.

In the year after the accident, the Nuclear Regulatory Commission (NRC) introduced its TMI Action Plan, which emphasized reducing operator error through control room redesign and increased training (U.S. NRC, 1980a, 1980b, 1981). The TMI accident also prompted the NRC to monitor plant operations more closely, increasing both plant inspections and the requisite detail in utility reporting of abnormal operating events (Baer, 1992; Jordan et al., 1988; Murley, 1990; Feinstein, 1989). In the late 1980s the NRC's expanded monitoring led to unprecedented multiyear shutdowns in several plants that the agency deemed to be operating unsafely (Warrock, 1987; Yates, 1988; National Research Council, 1992).

The nuclear utilities responded to the TMI accident by creating the Institute of Nuclear Power Operations (INPO), a self-monitoring body, charged with detecting poor nuclear plant performance and assisting member utilities in eliminating deficiencies (Pate, 1988; Moore, 1989). By the end of 1980, all U.S. utilities operating or building nuclear power plants were members of INPO (Thomas, 1988). By 1985, INPO had more than 400 employees and an annual budget of $44 million (Cook, 1986). INPO's major programs include: (1) dissemination of information on operating experience at member utility plants, (2) regular on-site plant evaluations, (3) accreditation of reactor personnel training programs, and (4) technical assistance for correction of recurring equipment and maintenance problems (Pate, 1988; Moore, 1989).

NRC directives and the activities of INPO in the years following TMI had a particularly dramatic impact on the training of plant operators and maintenance staff. Between 1979 and 1989, training personnel in the U.S. nuclear industry increased eightfold, and floor space devoted to training increased more than sixfold (Moore, 1989). In 1979, U.S. nuclear utilities had only 12 full-scale control room simulators, but, by 1988, this number had increased to 71 (Moore, 1989; Gonsalves, 1987). These increases in resources devoted to training far outpaced the 66% increase in the number of operating U.S. nuclear power plants during the decade after the TMI accident.

The question addressed in this study is whether the efforts of the nuclear utilities and the NRC to improve plant safety after TMI were accompanied by significant changes in performance records. As a step toward providing an answer, we examine temporal changes in a single safety-related performance measure: the hazard rate for unplanned outages. Our intertemporal comparisons are based on a Cox (1972, 1975) proportional hazards model estimated using operating spell duration data from 47 U.S. nuclear power plants. The 47 plants include all U.S. nuclear plants with net generating capacity of over 200 megawatts (MWe) that were in commercial operation prior to January 1, 1976. All these plants were operating for more than three years prior to the TMI accident, and for more than four years prior to any of the major regulatory reforms prompted by the accident. This set of plants thus provides us with a sufficient baseline sample of operating experience to facilitate useful comparisons with post-reform performance.

The parameters of the estimated model provide a measure of the relative risk of unplanned outages for each plant in each of three successive periods: 1976–1979, 1980–1985, and 1986–1991. The first of these periods ends in the year of the TMI accident, and precedes both the introduction of the NRC's TMI-Action Plan and the formation of INPO. The final period ends more than 12 years after the accident, and thus provides substantial leeway for detecting the effects of responses to TMI that entailed multiyear lags. We refer to the three periods as Pre-TMI, TMI-Transition, and Post-TMI.

All light water reactors undergo planned outages for routine maintenance and refueling. Prior to a planned outage, the plant is generally operating normally, and the outage begins with a gradual, manual shutdown. In contrast, unplanned outages are nearly always initiated by indications of malfunctions, most often due to equipment failures or operator errors. These unplanned outages either begin with a rapid shutdown of the reactor, known as a "scram," or are caused by conditions that would likely result in a scram if the plant were not shut down. The hazard rate for unplanned outages is of interest as an indicator of plant safety. This is because the malfunctions that trigger these outages may be complicated by additional failures caused by the stresses of a reactor scram.

The term "scram" refers specifically to the rapid, automated insertion of the control rods into the reactor core to halt fission and prevent overheating. Scrams that occur while the facility is generating moderate to high levels of power subject plant equipment to extreme mechanical and thermal loads. Such scrams increase the risk of a serious accident, since they could precipitate additional equipment failures, compounding the malfunction that originally triggered the shutdown (Jordan et al., 1988; Wiegle, 1986; Dredemis and Fourest, 1987; Denton, 1987). Further, the safety systems that trigger most scrams have their own probabilities of failure. A higher rate of scrams implies more trials of these safety mechanisms, and a higher likelihood that these safeguards will fail, allowing the reactor to continue operating under hazardous conditions (Jordan et al., 1988; Dredemis and Fourest, 1987). Apart from the stress which they place on mechanical components, frequent scrams may increase the likelihood of human error. If scrams become routine, operators may fail to follow established emergency procedures. This, in turn, may lead to inappropriate responses to serious malfunctions when they arise (Perrow, 1984; Dredemis and Fourest, 1987).

Nuclear power plants have layers of safety systems that are designed to mitigate occurrences that threaten the integrity of the reactor's major components. Multiple human and/or mechanical failures are required to penetrate these layers of protection. Because scrams elevate the risk of multiple failures, they hold key positions in the event trees used to estimate the likelihood of severe accidents in probabilistic risk analyses (PRA) of nuclear plants (U.S. NRC, 1975, 1989). The safety relevance of scrams is recognized by the NRC and INPO. In the years since TMI, both organizations have explicitly emphasized reducing scram frequencies through improved plant maintenance practices and operator training (Jordan et al., 1988; Murley, 1990; National Research Council, 1992; Moore, 1989; Denton, 1987).

Rates of scrams and unplanned outages during particular calendar time periods are routinely calculated by INPO and the NRC, and commonly appear in nuclear industry journals (Morimoto, 1986; Cletcher, 1989). Figure 1 contains examples of these simple performance measures for the 47 plants that we study in each of the three time periods defined above. These statistics clearly suggest a substantial decline from Pre-TMI levels in the number of unplanned outages and unplanned scrams per plant year, and per year of reactor operation. For three major reasons, however, the meaning and significance of these statistics is not entirely clear.

First, when rates of unplanned outages are calculated on a per plant basis, no adjustment is made for differences in the portion of the comparison periods during which the sample plants were at risk of outage. Thus, at least some of the evident decline in unplanned outages per plant year in figure 1 owes to an increase in the amount of time that the sample plants were out of service during the years after the TMI accident. Such declines in the annual rate of unplanned outages per plant need not be associated with any change in the plants' risk of such outages while running.

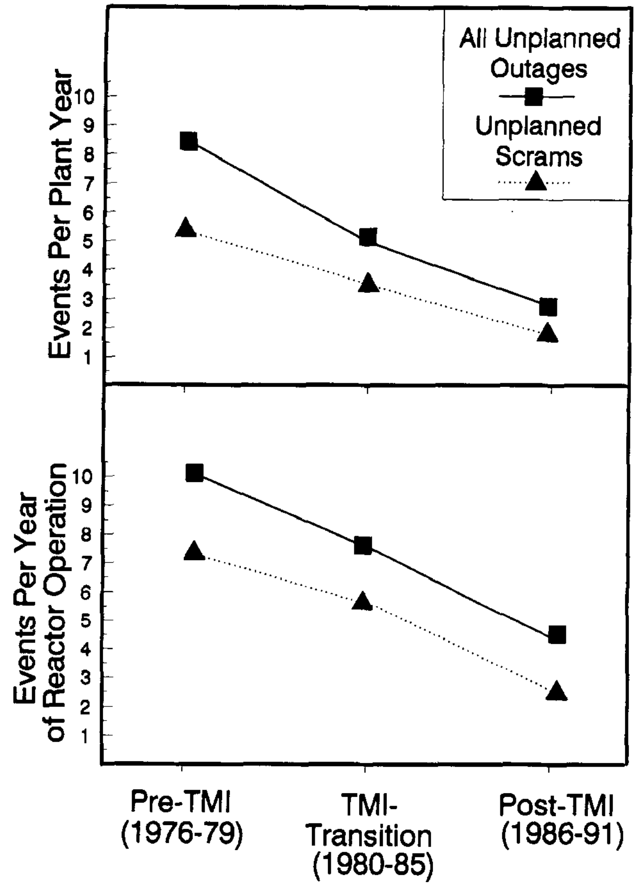

Figure 1. Unplanned outages and unplanned scrams per plant year and per year of reactor operation in the 47 sample plants during the Pre-TMI, TMI-Transition, and Post-TMI periods.

Second, while computing the rate of unplanned shutdowns per year of reactor operation is an improvement on crude "per plant year" rates, this adjustment for differences in exposure to the risk of unplanned outages assumes that the outage hazard is constant within and across spells of operation. Such statistics do not adjust for systematic variations in unplanned outage risk that depend on the timing of the operating spell in relation to major maintenance and repair activities. The hazard of unplanned outage is generally greater during an operating spell that follows a long outage for extensive repairs, than during a spell following a brief shutdown for tests or minor adjustments.[1] Further, prior research (Rust and Rothwell, 1996; Sturm, 1991) indicates that the unplanned outage hazard in nuclear plants tend to fall rapidly in the initial days of an operating spell, remain relatively stable over moderate spell durations, and then rise again at long spell durations. Performance comparisons across individual plants or across the same plant in different time periods are likely to be sensitive to a lack of adjustment for such systematic variations in unplanned outage risk. This is primarily because such plant specific comparisons must typically be based on relatively small samples of operating spells.

Finally, industry-wide unplanned outage rates may be misleading indicators of performance at the individual plant level, given that U.S. reactors are characterized by considerable heterogeneity. The potentially global consequences of an accident in a single nuclear plant imply that it is the individual plant hazards that are most relevant to assessing improvement in safety after TMI. The impression of improved safety given by figure 1 would be diminished if the declines in unplanned outages were, for example, known to be concentrated entirely in plants that had moderate to low rates of these events prior to TMI.

Our comparisons of unplanned outage risk using a Cox proportional hazards model address and rectify these deficiencies in the standard outage rate statistics that have come into vogue in the U.S. nuclear industry. They produce a substantially modified picture of the way in which these risks evolved in the aftermath of TMI. Specifically, our results suggest that most of the decline in unplanned outage risk occurred *after* the TMI-Transition period. This contrasts with the impression given by figure 1, which shows reductions in event frequencies between the Pre-TMI and TMI-Transition periods that are roughly comparable to those between the TMI-Transition and Post-TMI periods. The results also indicate that the decline in unplanned outage risk was concentrated in plants that had relatively high Pre-TMI unplanned outage risks. Most plants with moderate or low Pre-TMI unplanned outage risks experienced no significant risk reduction.

Previous econometric studies of nuclear power plant performance (Joskow and Rozanski, 1979; Easterling, 1982; Rothwell, 1990; Krautman and Solow, 1992; Lester and McCabe, 1993) have focused on the analysis of plant availability and capacity factors. These are primarily measures of productivity, more appropriate for gauging a plant's economic performance than its safety record. Many unplanned outages last less than a day, and thus large changes in the unplanned outage hazard can occur without reflecting themselves commensurately in capacity or availability factor movements. Alternative statistical approaches to the issue of U.S. nuclear plant safety can be found in Feinstein (1989), who studies detection of safety violations using NRC plant inspection records; in Dubin and Rothwell (1989), who model utility decisions on the timing of new safety system installations; and in PRA studies of severe accident risks, such as those undertaken by the NRC (U.S. NRC, 1975, 1989).

1. Description of the duration data

We obtained the spell duration data by calculating hours between successive periods of outage reported in the NRC's *Licensed Operating Reactors—Status Summary Report*. The sample consists of 5349 operating spells, and includes all spells from the 47 plants that began within the 16-year sample period. The analysis is conditioned on spells *beginning* in the sample period, to avoid dealing formally with left censoring. All spells are followed until the occurrence of an outage (planned or unplanned).

Table 1 provides descriptive statistics for the sample of operating spells in each of the three comparison periods, and for all periods combined. Equipment failures and operator errors initiate nearly 80% of the unplanned outages in our sample. The category "other unplanned outages" consists primarily of shutdowns following abnormal operating events or "transients" that trigger a scram, but are not subsequently traced to equipment failures or operator errors.

More than 70% of unplanned outages begin with scrams. While unplanned outages initiated with gradual, manual shutdowns have a less direct connection to plant safety, the frequency of these events is nonetheless related to the risk of serious accident. The same underlying malfunctions that cause these less precipitous shutdowns can also initiate unplanned automatic scrams, and would in most cases lead to a scram if the plant were not shutdown by other means.

Nearly a third of our sample of operating spells ends in a planned outage. Refueling and routine maintenance account for more than 80% of planned shutdowns. The "other planned outages" category includes outages for operator training and licensing tests, outages imposed by regulators, and outages undertaken for administrative reasons. Notably, more than 20% of planned outages are initiated by reactor scrams. During planned shutdowns, reactor scrams usually occur while the plant is producing low levels of power, or while the turbine generator is off-line. Such scrams do not produce the extreme stresses on plant equipment associated with unplanned scrams. Further, before a planned outage, the plant is typically operating normally, which diminishes the chance that a scram will

Table 1. Descriptive statistics for the reactor operating spell duration data

	Pre-TMI (1/1/76–12/31/79)	**TMI transition** (1/1/80–12/31/85)	**Post-TMI** (1/1/86–12/31/91)	**All periods combined** (1/1/76–12/31/91)
N	1967	2081	1301	5349
Mean Duration (days)	26.1	32.9	53.3	35.4
Standard Deviation (days)	33.1	45.3	70.8	50.3
Minimum Duration (days)	<0.1	<0.1	<0.1	<0.1
Maximum Duration (days)	392.7	439.9	480.4	480.4
Distribution of Outages of by Type (% of all outages):				
Planned Outages	27.4%	32.6%	39.0%	32.2%
Refueling	6.8%	8.0%	11.3%	8.4%
Maintenance	16.2%	18.6%	20.3%	18.1%
All Other	4.3%	6.0%	7.4%	5.7%
Unplanned Outages	71.6%	67.4%	61.0%	67.8%
Equipment Failure	59.3%	51.5%	47.2%	53.4%
Operator Error	9.9%	7.1%	4.8%	7.6%
All Other	3.4%	8.8%	9.0%	6.8%
Outages Initiated by Scrams (% of each type):				
Planned Outages	15.6%	32.1%	21.9%	24.0%
Unplanned Outages	71.7%	74.6%	62.2%	70.8%

lead to a coincidence of multiple failures. For these reasons, planned scrams lack the safety significance of unplanned scrams. In the hazard model estimation, all planned outage times, regardless of the method of shutdown, are treated as independent censoring times. Thus, the 32% of our sample spells that end in planned outages are treated as right censored in the hazard model estimation.

2. Methods

2.1 Proportional hazards model

In the context of modeling reactor operating spell durations, a hazard function describes the instantaneous risk of an outage at duration t, given that a spell lasts at least until t. The hazard model which we use assumes that if operating spell i, $i = 1 \ldots N$, occurs in plant k, $k = 1 \ldots K$, during period j, j = Pre-TMI, TMI-Transition, Post-TMI, then the hazard of unplanned outage for this spell is given by

$$\lambda_k^j(t; x_i) = \lambda_0(t)\exp[\alpha_k^j + x_i'\beta], \tag{1}$$

where

t is time elapsed from the start of the spell;
$\lambda_0(t)$ is an arbitrary and unspecified baseline hazard function;
α_k^j is a plant- and period-specific fixed effect applying to spells in plant k that begin in period j;
x_i is a vector of other observed characteristics for spell i;
and β is a conforming vector of unknown parameters to be estimated.

A hazard model of the type given in (1), which applies to durations between multiple failures, is sometimes called a modulated renewal process model (Kalbfleisch and Prentice, 1980). The hazard of unplanned outage is reset at the start of each new operating spell, but the hazard function varies across spells with the plant and period fixed effects, and other characteristics that are fixed at the start of the spell in question.

The vector of spell characteristics x_i includes functions of the time elapsed from the end of the last refueling outage in the plant to the start of spell i, as well as functions of the duration of the outage preceding the spell. Specifically, x_i includes the first three integer powers of time from last refueling to the start of the spell. And, with respect to the length of the outage preceding the spell, x_i is specified such that $\ln\lambda_k^j(t; x)$ is a linear spline function of this variable. The spline has two segments: one for outages less than ten days, and a second applying over outages of ten days or more. More general specifications were also considered, including one making $\ln\lambda_k^j(t; x)$ a cubic spline function of both time since refueling and length of last outage. These more flexible specifications yield no significant improvement in the fit of the estimated model, and give nearly identical estimates of the fixed effects.

This specification of x_i is used to account for changes in unplanned outage risk associated with the position of an operating spell relative to outages containing intensive plant maintenance and repair activities. The analysis treats such variations in outage risk as a common feature of light water reactor technology that is present in all sample plants. Functions of time from last refueling are included in x_i, because most major repair and maintenance operations in nuclear plants are done during the long outages required for refueling, which occur at intervals of 12 to 24 months. Between refuelings, equipment problems are most often minimally rectified in order to restore the plant to operation quickly, minimizing the need to purchase replacement power. The extent of repair and maintenance activities conducted during an outage is likely to be correlated with the length of the outage, which motivates the inclusion of functions of this variable in x_i. Our modulated renewal model is consistent with the results of Sturm (1991), whose tests based on operating spell duration data for European light water reactors rejected a simple, semi-Markov renewal model, but suggested the validity of treating the unplanned outage hazard as reset at the end of every refueling outage.

The spell characteristics x_i also include a binary indicator for whether the spell begins within a plant's first year of commercial operation. Eight of the 47 plants which we study came on-line less than a year before the start of the 16-year sample period. New nuclear power plants typically undergo a "teething" or "shakedown" period characterized by a high rate of outages for testing and adjustment of new plant equipment. Performance during this initial period of operation is often atypical of a plant's subsequent rate of unplanned outages. The first-year-of-operation indicator is included in x_i to prevent the estimated changes in the unplanned outage risk across periods in the newer sample plants from being dominated by these "shakedown" effects.

The functions of time since last refueling that are included in x_i could be more naturally introduced into the model as time-dependent covariates, so that time from refueling would accumulate over the operating spell, as opposed to being fixed at the start of the spell. However, with 5349 observations and 145 model covariates, allowing time-dependent covariates poses an intractable computational burden. Estimates of time-dependent and fixed covariate versions of an analogous hazard model for several smaller subsets of sample plants were nearly identical across specifications. Evidently, treating time from refueling as a time-dependent, rather than a fixed, covariate would have little impact on the estimated model.

2.2. Estimation

Estimates were obtained by maximization of the Breslow (1974) partial likelihood function with respect to $\gamma = (\alpha', \beta')'$, where α is the vector of plant and period-specific fixed effects. The estimated covariance used for inference is $\hat{V}(\hat{\gamma}) = -[\partial^2 \ln L(\hat{\gamma})/\partial\gamma^2]^{-1}$, where $\hat{\gamma}$ is the maximum likelihood estimate of γ, and $L(\gamma)$ is the partial likelihood function. We used the SAS, Version 6, proportional hazards regression procedure (PHREG) to compute the parameter estimates and the associated covariance matrix.

2.3. *Measures of unplanned outage risk*

We use the parameter estimates to construct measures of the hazard of unplanned outage as a proportion of the geometric mean hazard for all plants in the Pre-TMI period. Let $\overline{\lambda^{\text{Pre}}(t; x)} = \Pi_k \lambda_k^{\text{Pre}}(t; x)^{\frac{1}{K}}$ denote the geometric mean of the hazard function over all K plants for a spell with characteristics x that begins in the Pre-TMI period. We refer to $\overline{\lambda^{\text{Pre}}(t;x)}$ as the Pre-TMI "mean risk" of unplanned outage. Mean risk measures for the TMI-Transition and Post-TMI periods are analogously defined. Using this notation, the five unplanned outage risk measures we focus on are as follows:

$\theta_k^{\text{Pre}} = \lambda_k^{\text{Pre}}(t; x)/\overline{\lambda^{\text{Pre}}(t; x)} = \exp[\alpha_k^{\text{Pre}} - \frac{1}{K}\sum_h \alpha_h^{\text{Pre}}]$ is the Pre-TMI risk of unplanned outage in plant k as a proportion of the Pre-TMI mean risk;

$\theta_k^{\text{Transition}} = \lambda_k^{\text{Transition}}(t; x)/\overline{\lambda^{\text{Pre}}(t; x)} = \exp[\alpha_k^{\text{Transition}} - \frac{1}{K}\sum_h \alpha_h^{\text{Pre}}]$ is the TMI-Transition risk of unplanned outage in plant k as a proportion of the Pre-TMI mean risk;

$\theta_k^{\text{Post}} = \lambda_k^{\text{Post}}(t; x)/\overline{\lambda^{\text{Pre}}(t; x)} = \exp[\alpha_k^{\text{Post}} - \frac{1}{K}\sum_h \alpha_h^{\text{Pre}}]$ is the Post-TMI risk of unplanned outage in plant k as a proportion of the Pre-TMI mean risk;

$\theta^{\text{Transition}} = \overline{\lambda^{\text{Transition}}(t; x)}/\overline{\lambda^{\text{Pre}}(t; x)} = \exp[\frac{1}{K}\sum_k (\alpha_k^{\text{Transition}} - \alpha_k^{\text{Pre}})]$ is the TMI-Transition mean risk of unplanned outage as a proportion of the Pre-TMI mean risk; and

$\theta^{\text{Post}} = \overline{\lambda^{\text{Post}}(t; x)}/\overline{\lambda^{\text{Pre}}(t; x)} = \exp[\frac{1}{K}\sum_k (\alpha_k^{\text{Post}} - \alpha_k^{\text{Pre}})]$ is the Post-TMI mean risk of unplanned outage as a proportion of the Pre-TMI mean risk.

Estimates of θ_k^j and θ^j, which we write $\hat{\theta}_k^j$ and θ^j, are obtained by replacing the unknown effects by their maximum likelihood estimates. The 95% confidence interval bounds for θ_k^j are calculated as $\exp[\ln\hat{\theta}_k^j \pm 1.96\hat{V}(\ln\hat{\theta}_k^j)^{1/2}]$, where $\hat{V}(\ln\hat{\theta}_k^j) = h_k^{j\prime}\hat{V}(\hat{\gamma})h_k^j$, with h_k^j being a vector of constants such that $\ln\hat{\theta}_k^j = h_k^{j\prime}\hat{\gamma}$. Confidence intervals for θ^j are constructed in analogous fashion.

For our purposes, these relative risk measures have two major advantages. First, the unplanned outage risk in each plant is gauged against a measure of the central tendency of unplanned outage risks for comparable spells in all plants during the base period. This comparison is of greater intrinsic interest than the comparison to the risk in an arbitrarily chosen plant that is implicit in the estimated fixed-effects vector. Second, the measures are convenient for relating changes in each plant's unplanned outage risk to its Pre-TMI performance. In assessing efforts to improve performance after TMI, a question of central interest is whether plants with high Pre-TMI unplanned outage risk had any significant reduction in this risk after TMI.

2.4. *Hypothesis tests*

We use the estimated hazard model to test hypotheses relating to the changes in the sample plants' risks of unplanned outage across periods, and to differences in unplanned outage risk across plants. Table 2 provides formal statements of the hypotheses along with their interpretations. All of these hypotheses can be written in the form H_0: $L\gamma = c$, where L is a matrix of coefficients for the linear hypotheses, c is a vector of constants, and γ is the vector of parameters which we estimate. For testing these hypotheses, we use Wald chi-square statistics, which can be written in the form $\chi^2_w = (L\hat{\gamma} - c)'[L\hat{V}(\hat{\gamma})L']^{-1}(L\hat{\gamma} - c)$.

2.5. *Specification analysis*

We checked the appropriateness of the proportional hazards specification in (1) using the log-integrated hazard plotting procedure that is described by Anderson (1982), and that is also demonstrated in the text by Kalbfleisch and Prentice (1980). The procedure involves estimating a more general model, allowing for distinct, arbitrary baseline hazards within strata. This estimated model is then used to obtain log-integrated hazard function estimates for each stratum, and these are plotted against spell duration. We made two sets of these plots, one with strata corresponding to the three time periods of interest, and a second set with strata corresponding to the 47 sample plants.[2] The overall impression given by these log-integrated hazard plots is that the proportional hazards specification in (1) provides a useful approximation to the process generating the sample spells. In most instances the differences in the plots across strata approximate the constant separation that is implied by the proportional hazards specification in (1). For the plant strata case, there are instances in which the log-integrated hazard plots for different plants cross, but it is almost always in cases where the plots have similar shapes and lie nearly on top of one another.

Table 2. Hypotheses tested using the estimated fixed effects from the proportional hazards model

Null hypothesis*	Interpretation
$H_0: \theta_k^r - \theta_k^s = 0 \Leftrightarrow \alpha_k^r = \alpha_k^s$	No change in unplanned outage risk in plant k across periods r and s.
$H_0: \theta_k^r - \theta_k^s = 0, \forall k \Leftrightarrow \alpha_k^r = \alpha_k^s, \forall k$	No change in unplanned outage risk in any plant across periods r and s.
$H_0: \theta^r = \theta^s \Leftrightarrow \frac{1}{K}\sum_k \alpha_k^r = \frac{1}{K}\sum_k \alpha_k^s$	No change in the mean unplanned outage risk across periods r and s.
$H_0: \theta_k^r = 1 \Leftrightarrow \alpha_k^r = \frac{1}{K}\sum_h \alpha_h^{Pre}$	No difference between the Pre-TMI mean unplanned outage risk and the unplanned outage risk in plant k during period r.
$H_0: \theta_k^r = \theta^r, \forall k \Leftrightarrow \alpha_k^r = 0, \forall k$	No difference in unplanned outage risk across plants in period r.

*Each null is tested against the two-sided alternative obtained by changing the equalities to $\neq$.

We also used the plant-specific baseline hazards model to check the sensitivity of tests for change in outage risk across periods to relaxation of the common baseline hazard assumption. In this model, the hazard function for spells in plant k is given by $\lambda_{0k}(t)\exp[d_i^{\text{Transition}}\alpha_k^{\text{Transition}} + d_i^{\text{Post}}\alpha_k^{\text{Post}} + x_i'\beta]$, where $\lambda_{0k}(t)$ is the plant-specific baseline hazard, and d_i^j is a binary variable that takes a value of one when spell i begins in period j. Estimates of this model do not permit tests of the homogeneity of outage hazards across plants. The model can, however, be used to test hypotheses on changes in outage risk that are analogous to those in table 2. The results of these tests were nearly identical to the results presented below. Further, estimates of $\theta^{\text{Transition}}$ and θ^{Post} from this model were also nearly identical to those presented below.

Finally, we supplemented these informal checks of the proportional hazards specification in (1) with the generalized moments (GM) specification test of Horowitz and Neumann (1992). The GM test is based on the asymptotic distribution, under the proportional hazards assumption, of sample moments of functions of the generalized residuals (Horowitz and Neumann, 1992; Tsiatis, 1981). Unlike the Schoenfeld (1980) and Moreau, O'Quigley, and Mesbah (1985) specification tests, the GM test does not require an arbitrary partitioning of the data. Additionally, Monte Carlo experiments conducted by Horowitz and Neumann (1992) suggest that their GM test has substantially greater power against accelerated failure time alternatives than the Schoenfeld (1980) and Moreau, O'Quigley, and Mesbah (1985) tests. The GM test fails to reject the proportional hazards model in (1) at conventional significance levels ($\chi^2_{(1)} = 0.02, p = 0.90$).

3. Results

The proportional hazards model estimates show evidence of substantial change in the risk of unplanned outage across periods. The hypothesis of no change in unplanned outage hazard in any plant is strongly rejected by the Wald test for comparisons across each pair of periods ($p < 0.0001$). The hypothesis of no change in the mean unplanned outage risk across periods is also strongly rejected for each pair of periods ($p < 0.0001$). Estimates of the mean unplanned outage risk as a proportion of the Pre-TMI mean risk indicate a 20% decline in mean risk from the Pre-TMI period to the TMI-Transition period ($\theta^{\text{Transition}} = 0.80$, 95% C.I. = [0.73, 0.87]), and a nearly 50% decline in mean unplanned outage risk from the Pre-TMI period to the Post-TMI period ($\theta^{\text{Post}} = 0.51$, 95% C.I. = [0.46, 0.56]). The point estimates of $\theta^{\text{Transition}}$ and θ^{Post} thus suggest an acceleration in the rate of decline in unplanned outage risk after the TMI-Transition period. The estimated reduction in the mean risk between the TMI-Transition and Post-TMI periods is 1.5 times larger than that between the Pre-TMI and TMI-Transition periods.

Table 3 contains the estimated Pre-TMI ($\hat{\theta}_k^{\text{Pre}}$) and Post-TMI ($\hat{\theta}_k^{\text{Post}}$) relative risks of unplanned outage for each sample plant. Plants appear in table 3 in descending order of their estimated Pre-TMI unplanned outage risks. The table also contains the differences in the estimated relative risk measures for each plant across each pair of periods. These differences provide estimates of the changes in unplanned outage risk as a proportion of

Table 3. Relative risks of unplanned outage and changes in relative risk across periods

Pre-TMI risk rank	Plant, (MWe, reactor vendor, 1st year on line), utility, state	$\hat{\theta}_k^{Pre}$	$\hat{\theta}_k^{Post}$	$(\hat{\theta}_k^{Transition} - \hat{\theta}_k^{Pre})$	$(\hat{\theta}_k^{Post} - \hat{\theta}_k^{Transition})$	$(\hat{\theta}_k^{Post} - \hat{\theta}_k^{Pre})$
1.	Palisades, (812, CE, 1971)	2.13	0.89	−1.11*	−0.14	−1.25**
	Consumers Power, MI	[1.62, 2.79]	[0.590, 1.31]			
2.	Brunswick 2, (867, GE, 1975)	2.03	0.60	−0.44	−0.99**	−1.43**
	Carolina Power & Light, NC	[1.57, 2.64]	[0.38, 0.95]			
3.	Indian Point 2, (1013, W, 1974)	1.90	0.78	−0.61	$-0.51^{\dagger}$	−1.12**
	Consolidated Edison, NY	[1.42, 2.56]	[0.53, 1.15]			
4.	Browns Ferry 1, (1152, GE, 1974)	1.83	—	−0.53	—	—
	Tennessee Valley Authority, AL	[1.36, 2.47]				
5.	Hatch 1, (850, GE, 1975)	1.80	0.66	0.03	−1.17**	−1.14**
	Georgia Power, GA	[1.38, 2.33]	[0.43, 0.99]			
6.	Browns Ferry 2, (1152, GE, 1975)	1.78	0.69	−0.28	−0.81	−1.08
	Tennessee Valley Authority, AL	[1.31, 2.40]	[0.26, 1.84]			
7.	Oconee 1, (887, BW, 1973)	1.68	0.50	$-0.60^{\dagger}$	−0.59*	−1.18**
	Duke Power, SC	[1.24, 2.26]	[0.31, 0.78]			
8.	Zion 1, (1098, W, 1973)	1.66	0.63	−0.55	$-0.48^{\dagger}$	−1.03**
	Commonwealth Edison, IL	[1.24, 2.22]	[0.40, 0.99]			
9.	Zion 2, (1098, W, 1974)	1.54	0.65	−0.35	$-0.53^{\dagger}$	−0.89*
	Commonwealth Edison, IL	[1.12, 2.10]	[0.43, 0.98]			
10.	Robinson 2, (769, W, 1971)	1.42	0.71	−0.01	−0.71*	−0.72*
	Carolina Light & Power, SC	[1.06, 1.92]	[0.48, 1.04]			
11.	Millstone 2, (910, CE, 1975)	1.33	0.60	$-0.55^{\dagger}$	−0.18	−0.73*
	Northeast Nuclear Energy, CT	[0.97, 1.81]	[0.39, 0.91]			
12.	Dresden 3, (810, GE, 1971)	1.32	0.77	−0.05	$-0.51^{\dagger}$	$-0.55^{\dagger}$
	Commonwealth Edison, IL	[0.96, 1.83]	[0.51, 1.16]			
13.	Turkey Point 4, (760, W, 1973)	1.31	1.17	−0.01	−0.13	−0.14
	Florida Power & Light, FL	[0.95, 1.80]	[0.82, 1.67]			

Notes: Bracketed figures are 95% confidence interval bounds. Changes in risk significantly different from zero in the Wald tests are indicated by † at the 5% significance level, by * at the 1% level, and by ** at the 0.1% level. For the Three Mile Island 1 plant we report only the estimated change in risk across the Pre- and Post-TMI periods. This is because the TMI-1 plant operated for only two months during the TMI-Transition period. See text for definitions of the $\hat{\theta}_k^j$, j = Pre, Transition, Post. $\hat{\theta}_k^{Post}$ not estimated for Browns Ferry 1, because the plant had no operating spells in the Post-TMI period. Reactor vendors: CE = Combustion Engineering, GE = General Electric, BW = Babcock and Wilcox, W = Westinghouse.

Table 3 (continued). Relative risks of unplanned outage and changes in relative risk across periods

Pre-TMI risk rank	Plant, (MWe, reactor vendor, 1st year on line), utility, state	$\hat{\theta}_k^{Pre}$	$\hat{\theta}_k^{Post}$	$(\hat{\theta}_k^{Transition} - \hat{\theta}_k^{Pre})$	$(\hat{\theta}_k^{Post} - \hat{\theta}_k^{Transition})$	$(\hat{\theta}_k^{Post} - \hat{\theta}_k^{Pre})$
14.	Calvert Cliffs 1, (918, CE, 1975)	1.31	0.60	−0.22	−0.48†	−0.70*
	Baltimore Gas & Electric, MD	[0.96, 1.79]	[0.38, 0.96]			
15.	Oconee 2, (887, BW, 1974)	1.29	0.42	−0.43	−0.44†	−0.87**
	Duke Power, SC	[0.92, 1.81]	[0.25, 0.69]			
16.	Peach Bottom 3, (1152, GE, 1974)	1.27	0.95	−0.59†	0.27	−0.31
	Philadelphia Electric, PA	[0.91, 1.77]	[0.61, 1.49]			
17.	Oconee 3, (887, BW, 1974)	1.24	0.52	−0.21	−0.51†	−0.72*
	Duke Power, SC	[0.88, 1.74]	[0.33, 0.83]			
18.	Fitzpatrick, (883, GE, 1975)	1.21	0.47	−0.49	−0.26	−0.74*
	New York Power Authority, NY	[0.86, 1.71]	[0.29, 0.77]			
19.	Duane Arnold, (566, GE, 1975)	1.21	0.57	−0.28	−0.36	−0.64†
	Iowa Electric Power & Light, IA	[0.84, 1.74]	[0.37, 0.89]			
20.	Quad Cities 1, (828, GE, 1973)	1.19	0.65	−0.43	−0.11	−0.55†
	Commonwealth Edison, IL	[0.86, 1.65]	[0.43, 0.97]			
21.	Rancho Seco 1, (929, BW, 1975)	1.15	0.68	0.22	−0.69	−0.47
	Sacramento Municipal Utility, CA	[0.80, 1.65]	[0.29, 1.60]			
22.	Turkey Point 3, (760, W, 1972)	1.15	1.11	−0.05	0.01	−0.04
	Florida Power & Light, FL	[0.82, 1.60]	[0.75, 1.64]			
23.	Cook 1, (1089, W, 1975)	1.14	0.32	−0.62*	−0.20	−0.82**
	Indiana & Michigan Electric, MI	[0.82, 1.57]	[0.18, 0.59]			
24.	Surry 1, (848, W, 1972)	1.11	0.56	0.10	−0.65*	−0.55†
	Virginia Power, VA	[0.75, 1.63]	[0.34, 0.90]			
25.	Peach Bottom 2, (1152, GE, 1974)	1.11	0.96	−0.25	0.10	−0.14
	Philadelphia Electric, PA	[0.78, 1.58]	[0.64, 1.45]			
26.	Surry 2, (848, W, 1973)	1.10	0.37	0.15	−0.88**	−0.72*
	Virginia Power, VA	[0.72, 1.67]	[0.21, 0.66]			
27.	Dresden 2, (810, GE, 1970)	1.10	0.84	0.34	−0.60†	−0.25
	Commonwealth Edison, IL	[0.77, 1.55]	[0.58, 1.23]			
28.	Pilgrim 1, (655, GE, 1972)	1.10	0.51	−0.21	−0.37	−0.58†
	Boston Edison, MA	[0.75, 1.60]	[0.27, 0.95]			

Table 3 (continued). Relative risks of unplanned outage and changes in relative risk across periods

Pre-TMI risk rank	Plant, (MWe, reactor vendor, 1st year on line), utility, state	$\hat{\theta}_k^{Pre}$	$\hat{\theta}_k^{Post}$	$(\hat{\theta}_k^{Transition} - \hat{\theta}_k^{Pre})$	$(\hat{\theta}_k^{Post} - \hat{\theta}_k^{Transition})$	$(\hat{\theta}_k^{Post} - \hat{\theta}_k^{Pre})$
29.	Quad Cities 2, (828, GE, 1973)	1.04	0.58	−0.43	−0.03	−0.46†
	Commonwealth Edison, IL	[0.72, 1.51]	[0.38, 0.89]			
30.	Arkansas 1, (902, BW, 1974),	0.96	0.81	0.40	−0.54†	−0.15
	Arkansas Power & Light, AR	[0.63, 1.45]	[0.56, 1.19]			
31.	Kewaunee, (535, W, 1974)	0.90	0.27	−0.48†	−0.16	−0.63*
	Wisconsin Public Service, WI	[0.61, 1.32]	[0.14, 0.50]			
32.	Millstone 1, (662, GE, 1971)	0.84	0.35	−0.23	−0.26	−0.48†
	Northeast Nuclear Energy, CT	[0.57, 1.22]	[0.20, 0.62]			
33.	Monticello, (569, GE, 1971)	0.80	0.39	−0.48†	0.07	−0.41†
	Northern States Power, MN	[0.53, 1.20]	[0.23, 0.66]			
34.	San Onofre 1, (450, W, 1968)	0.80	0.51	−0.01	−0.29	−0.30
	Southern California Edison, CA	[0.52, 1.23]	[0.30, 84]			
35.	Prairie Island 2, (593, W, 1974)	0.79	0.24	−0.46†	−0.09	−0.55*
	Northern States Power, MN	[0.52, 1.19]	[0.12, 0.48]			
36.	Prairie Island 1, (593, W, 1973)	0.78	0.18	−0.47*	−0.13	−0.60**
	Northern States Power, MN	[0.52, 1.17]	[0.09, 0.37]			
37.	Ginna, (517, W, 1970)	0.76	0.43	−0.49*	0.16	−0.33
	Rochester Gas & Electric, NY	[0.50, 1.15]	[0.27, 0.70]			
38.	Maine Yankee, (830, CE, 1972)	0.72	0.61	0.19	−0.30	−0.12
	Maine Yankee Atomic Power, ME	[0.45, 1.16]	[0.40, 0.92]			
39.	Fort Calhoun 1, (481, CE, 1974)	0.72	0.31	−0.25	−0.16	−0.41†
	Omaha Public Power, NE	[0.46, 1.12]	[0.18, 0.55]			
40.	Haddam Neck, (600, W, 1968)	0.67	0.43	0.09	−0.32	−0.23
	Connecticut Yankee Atomic, CT	[0.43, 1.04]	[0.24, 0.78]			
41.	Vermont Yankee, (514, GE, 1972)	0.66	0.28	0.11	−0.43†	−0.32
	Vermont Yankee Nuclear, VT	[0.38, 0.97]	[0.16, 0.52]			

Table 3 (continued). Relative risks of unplanned outage and changes in relative risk across periods

Pre-TMI risk rank	Plant, (MWe, reactor vendor, 1st year on line), utility, state	$\hat{\theta}_k^{Pre}$	$\hat{\theta}_k^{Post}$	$(\hat{\theta}_k^{Transition} - \hat{\theta}_k^{Pre})$	$(\hat{\theta}_k^{Post} - \hat{\theta}_k^{Transition})$	$(\hat{\theta}_k^{Post} - \hat{\theta}_k^{Pre})$
42.	Oyster Creek 1, (550, GE, 1969)	0.59	0.71	0.02	0.10	0.12
	General Public Utilities, NJ	[0.36, 0.97]	[0.46, 1.08]			
43.	Cooper Station, (835, GE, 1974)	0.50	0.36	0.01	−0.15	−0.14
	Nebraska Public Power, NE	[0.30, 0.84]	[0.21, 0.63]			
44.	Nine Mile Point 1, (642, GE, 1969)	0.44	0.54	0.01	0.03	0.11
	Niagra Mohawk Power, NY	[0.23, 0.84]	[0.30, 0.98]			
45.	Point Beach 1, (524, W, 1970)	0.34	0.13	−0.15	−0.05	−0.20
	Wisconsin Electric Power, WI	[0.18, 0.62]	[0.06, 0.32]			
46.	Point Beach 2, (524, W, 1972)	0.27	0.16	−0.11	<0.01	−0.11
	Wisconsin Electric Power, WI	[0.14, 0.53]	[0.07, 0.35]			
47.	Three Mile Island 1, (871, BW, 1974)	0.25	0.37	—	—	0.12
	General Public Utilities, PA	[0.11, 0.59]	[0.22, 0.63]			

the Pre-TMI mean risk for all sample plants. Cross-period differences in unplanned outage risk that are significantly different from zero at the 5% level or lower in the Wald tests are marked in table 3. Plants with unplanned outage risks significantly different from the Pre-TMI mean risk under the Wald test can be identified by inspection of the 95% confidence intervals for θ_k^{Pre} and θ_k^{Post} in table 3. A confidence interval not containing a value of one implies that the plant's estimated unplanned outage risk in the corresponding period is significantly different from the Pre-TMI mean risk at the 5% level.

From the estimates reported in table 3, it can be seen that the reductions in the mean unplanned outage risk from the base period to the Post-TMI period reflect reductions in risk in the majority of sample plants. The point estimates of relative unplanned outage risk in the Post-TMI period are lower than those for the Pre-TMI period in all but three sample plants. In 27 of the 47 sample plants, the estimated decline in unplanned outage risk between the Pre- and Post-TMI periods is significantly different from zero at the 5% level or lower under the Wald test. No plants had statistically significant increases in unplanned outage risk across the Pre- and Post-TMI periods.

While a number of plants had substantial declines in unplanned outage risk between the Pre-TMI and TMI-Transition periods, careful inspection of the estimates in table 3 suggests that most of the risk reduction in the sample plants occurred after the Transition period. A majority of the 47 plants had larger estimated reductions in unplanned outage risk between the first two sample periods than between the last two periods. In several plants that had substantial and statistically significant reductions in unplanned outage risk, virtually all of the reduction is estimated to occur after the Transition period. More generally, the mean estimated reduction in unplanned outage risk for the sample plants between the TMI-Transition and Post-TMI periods is 1.8 times the mean reduction between the Pre-TMI and TMI-Transition periods.[3] Further, while 17 sample plants had statistically significant reductions in unplanned outage risk from the TMI-Transition to the Post-TMI period, only ten plants had significant reductions across the Pre-TMI and TMI-Transition periods, and one plant had a statistically significant increase in unplanned outage risk across these periods.

The largest estimated reductions in unplanned outage risk were concentrated among those plants that had the highest estimated Pre-TMI unplanned outage risks. This pattern is particularly clear in figure 2, where the Pre-TMI and Post-TMI relative risk estimates for each sample plant are plotted against the plants' Pre-TMI unplanned outage risk ranks. Among the plants with the 15 highest estimates of $\hat{\theta}_k^{Pre}$, 12 had statistically significant reductions in unplanned outage risk across the Pre-TMI and Post-TMI periods. The changes in risk among these plants appear quite substantial, with the eight highest Pre-TMI risk plants all having estimated absolute declines in unplanned outage risk exceeding the Pre-TMI mean risk level (i.e., $|\hat{\theta}_k^{Post} - \hat{\theta}_k^{Pre}| > 1$). In contrast, among the plants having the 15 lowest estimates of $\hat{\theta}_k^{Pre}$, only three had significant declines in unplanned outage risk over the sample period. Further, only five of these plants had estimated absolute reductions in unplanned outage risk across the Pre- and Post-TMI periods that exceeded 25% of the Pre-TMI mean risk level.

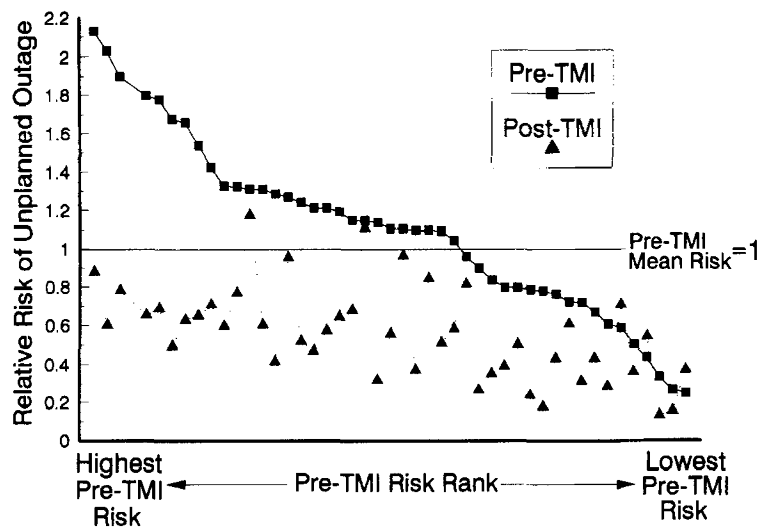

Figure 2. Estimated Pre-TMI ($\hat{\theta}_k^{Pre}$) and Post-TMI ($\hat{\theta}_k^{Post}$) relative risks of unplanned outage for individual plants plotted against the plants' Pre-TMI risk ranks. (Note: The plot excludes Browns Ferry 1, which had no spells of operation during the Post-TMI period.)

The reductions in the risk of unplanned outage in plants with relatively high Pre-TMI risk levels narrows the range of unplanned risk estimates. In the Pre-TMI period, the highest estimated unplanned outage risk for the sample plants is 2.13 (95% C.I. = [1.62, 2.79]) times the Pre-TMI mean risk, while the lowest risk estimate is only 0.25 (95% C.I. = [0.10, 0.59]) times the Pre-TMI mean risk. In the Post-TMI period, the range of relative risk estimates is nearly halved, with the highest value of $\hat{\theta}_k^{Post}$ being 1.17 (95% C.I. = [0.82, 1.67]), and the lowest value being 0.13 (95% C.I. = [0.06, 0.32]). Despite the reduced range, the highest unplanned outage risk estimate for the sample plants is more than eight times the lowest estimated risk in both the Pre- and Post-TMI periods. Further, the hypothesis of no difference in unplanned outage hazards across plants is rejected in each period by the Wald test ($p < 0.0001$).

As an additional means of depicting the pattern of change in unplanned outage risk we constructed survivor function plots (figure 3) based on estimates of a model with distinct, arbitrary baseline hazards in nine strata. The strata were defined by first partitioning plants into three groups based on the estimated fixed effects model: those with Pre-TMI outage risks significantly above the Pre-TMI mean under the Wald test (i.e., the first ten plants listed in table 3), those with Pre-TMI risk significantly below the Pre-TMI mean (i.e., the last seven plants listed in table 3), and all remaining plants. Spells occurring in each of the three plant groups were then subdivided by period (Pre-TMI, TMI-Transition, Post-TMI), yielding nine strata.

The hazard function for spell i from the qth stratum under this model is given by $\lambda_{0q}(t)\exp(x_i'\beta)$, where $\lambda_{0q}(t)$ is the arbitrary baseline hazard function, and $x_i'\beta$ is as defined before. The estimates plotted in figure 3 are of the form $\hat{S}_{0q}(t)^{\exp(\bar{x}'\hat{\beta})}$, where $\hat{S}_{0q}(t)$ is the

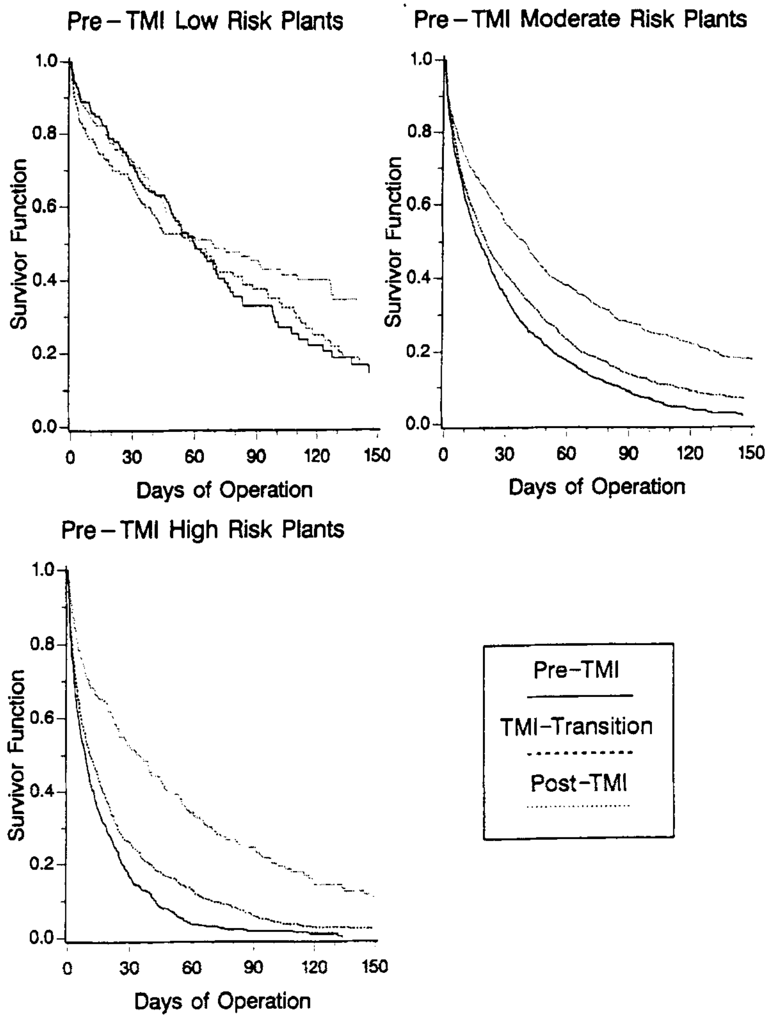

Figure 3. Plots of survivor function estimates derived from estimates of a proportional hazards model allowing distinct, arbitrary baseline hazards for strata defined by period and Pre-TMI unplanned outage risk.

nonparametric, maximum likelihood estimate of the baseline survivor function in the qth stratum (Kalbfleisch and Prentice, 1980), and $\hat{\beta}$ is the estimated vector of hazard model parameters. Here $\tilde{x}$ is used to denote the spell characteristics vector evaluated at the sample means of time-since-last-refueling and length-of-last-outage, and with the first year of operation indicator set to zero. Figure 3 illustrates the larger evident reduction in the probability of unplanned outage among plants with high Pre-TMI risk levels, and the concentration of most of the change in unplanned outage risk between the TMI-Transition and Post-TMI periods.

4. Discussion

Our proportional hazards model estimates clearly suggest that the sample plants experienced large reductions in unplanned outage risk during the years after TMI. How much of the change in this safety-related measure can be attributed to NRC and utility efforts to improve reactor safety after the TMI accident? Our results cannot provide a definitive answer to this question. However, two aspects of the pattern of risk reduction evident in our estimates make it plausible to believe that these safety reform efforts may have played a major role.

First, the fact that most of the change in unplanned outage risk appears to occur from the TMI-Transition period to the Post-TMI period is not consistent with what one would expect from nuclear plants moving along stable "learning curves." Nuclear plants tend to experience improvements in performance with age, as equipment and operating protocols are adjusted to optimize plant performance.[4] In the absence of major technological innovations, or other changes in the operating environment, the pace of performance improvement would generally be expected to diminish with age. Indeed, studies that have related nuclear power plant capacity factors to functions of years in operation (Joskow and Rozanski, 1979; Krautman and Solow, 1992; Lester and McCabe, 1993; Sturm, 1993) suggest concave relationships between plant performance and operating experience, with the pace of improvement slowing dramatically after the first decade of operation.

Our estimates suggest an acceleration in performance improvement that occurs from the TMI-Transition period to the Post-TMI period. By the time of this acceleration, the sample plants are either entering or well into their second decade of operation. Such an increase in the pace of performance improvement conflicts with the hypothesis of plants moving along stable, concave age-performance curves. It is, however, consistent with the timing of the reforms undertaken in response to TMI. These reforms included new NRC and INPO monitoring programs, as well as innovations in personnel training, control room design, plant instrumentation, and plant maintenance practices. The changes introduced in these areas were in many cases large scale undertakings that took years to complete. Thus, it is reasonable to expect that their effects would not be fully evident until the Post-TMI period.

A second pattern that suggests that the reforms after TMI may have played a major role in reducing unplanned outage risk levels is the evident concentration of these reductions among the plants with the highest Pre-TMI unplanned outage risks. Both the NRC and INPO made reduction of reactor scrams through improvements in plant maintenance and operator training an explicit goal after TMI (Jordan et al., 1988; Murley, 1990; National Research Council, 1992; Moore, 1989; Denton, 1987). It is reasonable to expect that a disproportionate share of these scram reduction efforts would be focused on "outlier" plants, with relatively high risks of unplanned outages. These plants were likely to be perceived by INPO and NRC regulators as posing relatively high accident risks. They may also have been viewed as having the potential for achieving dramatic reductions in scram frequency at relatively low cost, through adoption of maintenance and operating protocols already in use at plants with better performance records. Moreover, one must consider the impact of the NRC's aggressive posture in shutting down plants which it deemed to be operating unsafely during the Post-TMI period. This provided a strong incentive for the utilities operating these outlier plants to bring their unplanned outage risk levels into line with the rest of the industry.[5]

One of the more remarkable features of our results is the degree of variation in the estimated unplanned outage risk across sample plants. Previous studies of nuclear plant performance relate such variations to plant size, age, and reactor vendor (Joskow and Rozanski, 1979; Easterling, 1979; Rothwell, 1990; Krautman and Solow, 1992; Lester and McCabe, 1993; Sturm, 1993). Casual inspection of table 3 suggests some support for these factors as correlates of unplanned outage risk. For instance, smaller, older reactors are overrepresented among these plants with the lowest estimated unplanned outage risks in both the Pre- and Post-TMI periods. Still, our unplanned outage risk estimates clearly point to the limits of such generalizations. It is easy to find instances where plants of similar size and age with the same reactor vendor are at opposite ends of the spectrum of risk estimates. Compare, for example, the relative risk estimates in table 3 for Brunswick 2 and Cooper Station in the Pre-TMI period, or for Indian Point 2 and Cook 1 in the Post-TMI period.

One set of likely contributors to these variations in performance across ostensibly similar plants includes differences in the quality, skill, and organization of plant management in the U.S. nuclear industry (Marcus, 1988; Rothwell, 1995). A notable feature of our results that tends to support this management explanation is the similar performance of "sister" plants that are located on the same site and operated by the same utility. In ten of the 11 sets of such plants in our sample, there were no significant differences across sister facilities in either the level of unplanned outage risk, or the pattern of change in this risk across periods. In all but one case, these sister plants entered operation within two years of each other, and have reactors of the same size, built by the same vendor.[6] This probably contributes to their performance similarities. Still, there is evidently much greater uniformity in performance across reactors of similar design and age when they are operated by the *same* utility, than when they are operated by *different* utilities.

A second possible explanation for the wide variations in plant performance evident in our results is the limited degree of standardization in U.S. nuclear industry (Lester and McCabe, 1993; David and Rothwell, 1996). Even when two U.S. nuclear plants have reactors of the same design, manufactured by the same vendor, the remainder of the plants' facilities are likely to be designed by different architectural and engineering firms and built by different construction firms. Differences in the quality and reliability of these custom, "one-of-a-kind" installations may contribute to plants with similar reactors having very different performance.

The reductions in unplanned outage risk evident in our analysis are consistent with the hypothesis that U.S. nuclear plant safety has improved since the TMI accident. It must be kept in mind, however, that our analysis is restricted to a single, safety-related measure, and that no such measure can adequately subsume the complexities of nuclear plant safety. The limitations of the performance measure which we study here are perhaps best illustrated by the fact that the surviving sibling of the ill fated TMI-2 plant had the lowest unplanned outage risk in our sample during the Pre-TMI period (see table 3). The cooling system valve flaw that played a major role in the TMI-2 accident was also present in the

TMI-1 plant. However, this potentially dangerous problem in the TMI-1 plant did not manifest itself in a high rate of unplanned outages during the years prior to the TMI accident.

One possible criticism of unplanned outage risk as an indicator of nuclear plant safety is that operators may be able to reduce the frequency of unplanned outages by increasing the system's tolerance for abnormal operating events, rather than by altering plant maintenance, equipment, or operating procedures to lessen the frequency of such events. Under this scenario reduction in unplanned outage risk would not reflect a decline in the risk of potentially hazardous operating events, but an increase in the willingness to allow the plant to continue operating when such events occur.

This type of scenario may be applicable to reactors in countries with less restrictive regulatory regimes than the U.S., but it is implausible as an explanation for the decline in unplanned outage risk evident in our U.S. sample. The reactor safety systems required in U.S. nuclear plants have trip levels that cause the reactor to scram automatically in the presence of deviations from specified operating parameters. With the increased scrutiny of plant operations by the NRC in the wake of TMI, and with the NRC's 1984 introduction of mandatory reporting of all challenges to engineered safety features, it is difficult to believe that utilities would have had much leeway for reducing the stringency of these trip levels.

On the other hand, it is true that improved instrumentation adopted in a number of plants since TMI has in some cases allowed adjustment of trip levels to reflect the reduced margins of error in instrument reporting on the state of the plant. Even without a change in the underlying pattern of operating events, such improvements in instrumentation can reduce unplanned scrams by reducing instances of Type I error, in which a scram is unnecessarily initiated when the plant is in fact operating within the desired safety limits. But, since scrams are a source of accident risk per se, even reductions in Type I scrams represent a safety gain.

Given the potential global consequences of a severe nuclear accident—amply demonstrated by the case of Chernobyl—there are obvious social benefits associated with improvements in reactor safety. However, such safety gains are not in general without costs to taxpayers, utilities, and their customers. It is beyond the scope of this study to attempt a valuation of the relative costs and benefits associated with the changes in unplanned outage risk evident in our analyses. In considering such a calculation, however, it is important to note that reductions in unplanned outage risk, in addition to the safety related benefits they are likely to hold, can also lead to substantial improvements in plant productivity. A lower frequency of unplanned outages typically reduces plant operating losses from these outages, as well as the need for purchases of replacement power. In the 27 sample plants that had significant reductions in their unplanned outage risks, total operating losses from these outages fell by 125 days per year between the Pre-TMI and Post-TMI periods. At recent wholesale prices, this change in unplanned operating losses translates to a $48 million per year reduction in replacement power costs. In contrast, the 19 sample plants that experienced no significant change in their unplanned outage risks had increases in both their total unplanned operating losses and their associated replacement power costs over the sample period. The replacement power cost savings associated with lowering unplanned outage risk at least partially offsets the costs incurred in obtaining these risk reductions.

Acknowledgments

This research was supported in part by NSF Grant IRI-8814179-02/03 to the Center for Economic Policy Research Project on "Informational and Organizational Impacts on Productivity—The Economics of Control and Reliability in Complex Production Processes." The authors are grateful for comments on previous drafts of this article by Timothy Bresnahan, Roland Sturm, Ed Steinmueller, members of the Technology, Organizations and Productivity Workshop at Stanford University, and an anonymous referee.

Notes

1. In our data, this pattern is evident in comparisons of Kaplan–Meier survivor function plots across subsamples of operating spells that followed either short or long periods of outage. Sturm (1991) found similar patterns in his analyses of European light water reactors.
2. These plots are not provided here, but are available from the authors upon request.
3. The mean of the estimated changes in unplanned outage risk between the Pre-TMI and TMI-Transition periods, shown in Table 3, is −0.19. The mean of the changes in the plants' unplanned outage risks between the TMI-Transition and Post-TMI periods, also shown in table 3, is −0.34. The mean of the individual plants' estimated risk reductions between the latter two sample periods is thus 1.8 times larger than that between the first two periods.
4. Our results suggest considerable heterogeneity across plants in the pattern of change in unplanned outage risk with age. There is evidently a high degree of dispersion in both the plants'-initial unplanned outage risk levels and the rates of change—if any—in this risk over time. For instance, while many of the plants which we study experience substantial unplanned outage risk reductions, there are 19 sample plants, mostly with low to moderate Pre-TMI risk levels, that had no significant change in unplanned outage risk over the 16-year sample period. Some plants evidently attain relatively low levels of unplanned outage risk at a young age, and this level remains relatively stable over the remainder of the plant's life. Because of this heterogeneity, we elected not to include functions of plant age in our hazard model specification. Such a specification would arbitrarily assume the existence of some "normal" pattern of unplanned outage risk reduction that is an inherent feature of light water reactor technology. Our plant and period-specific fixed effects specification can be seen as approximating *plant-specific* age-performance curves with a distinct step function for each plant.
5. There is an alternative interpretation of the concentration of unplanned outage risk reduction in plants with high Pre-TMI risk levels that may occur to some readers. Exceptionally high or low values in a stochastic process tend to be followed by values closer to the mean—a phenomenon commonly called "regression to the mean." The large risk reductions in plants with high Pre-TMI risk levels could be interpreted as an expected regression to the mean. Under this explanation, the changes in unplanned outage risk in these high-risk plants would only reflect the level of random variation in the process generating the sample spells, and not a systematic shift in the affected plants' hazard functions. We view this interpretation as highly implausible, since it is inconsistent with the general pattern of our results. For instance, why is it that only the exceptionally high-risk plants regress toward the mean? In our result, the plants with exceptionally low Pre-TMI unplanned outage risks have remarkably stable performance over the 16-year sample period.
6. 6. The 11 sets of "sister" plants in our sample are: Browns Ferry 1 and 2; Oconee 1, 2, and 3; Zion 1 and 2; Turkey Point 3 and 4; Dresden 2 and 3; Peach Bottom 2 and 3; Quad Cities 1 and 2; Surry 1 and 2; Prairie

Island 1 and 2; Point Beach 1 and 2; and Millstone 1 and 2. The Millstone plants are the only set of sister plants that had significantly different unplanned outage risk performance. Notably, the Millstone plants were also the only set of sister plants that had different reactor vendors, and that came on-line more than two years apart.

References

Anderson, P. K. (1982). "Testing Goodness of Fit in Cox's Regression and Life Model," *Biometrics* 38, 67–77.

Baer, T. (1992). "Safety at a Glance: Upgrading the Displays in a Nuclear Plant Control Room," *Mechanical Engineering* 114, 56–60.

Breslow, N. E. (1974). "Covariance Analysis of Censored Survival Data," *Biometrics* 30, 89–99.

Cletcher, J. W. (1989). "Reactor Shutdown Experience," *Nuclear Safety* 30, 265–267.

Cook, J. (1986). "INPO's Race Against Time," *Forbes* 137(4), 54–58.

Cox, D. R. (1972). "Regression Models and Life-Tables (with Discussion)," *Journal of the Royal Statistical Society*, Series B, 34, 187–220.

Cox, D. R. (1975). "Partial Likelihood," *Biometrika* 62, 269–276.

National Research Council (1992). *Nuclear Power: Technical and Institutional Options for the Future*. Washington: National Academy of Sciences.

Cunningham, A. M. (1989). "Ten Years After: Cleaning Up Three Mile Island," *Technology Review* 92, 18–20.

David, P. A., and G. S. Rothwell. (1996). "Measuring Standardization: An Application to the American and French Nuclear Power Industries," *European Journal of Political Economy* (forthcoming).

Denton, H. R. (1987). "Reactor Scrams in the U.S.A.: A Regulators Point of View." In *Proceedings of an NEA Symposium on Reducing the Frequency of Reactor Scrams*. Paris: OECD.

Dredemis, G., and B. Fourest (1987). "Reactor Scrams, Why a Safety Concern?" In *Proceedings of an NEA Symposium on Reducing the Frequency of Reactor Scrams*. Paris: OECD.

Dubin, J. A., and G. S. Rothwell (1989). "Risk and Reactor Safety System Adoption," *Journal of Econometrics* 42, 201–18.

Easterling, R. G. (1982). "Statistical Analysis of U.S. Power Plant Capacity Factors Through 1979," *Energy* 7, 253–258.

Feinstein, J. S. (1989). "The Safety Regulation of U.S. Nuclear Power Plants: Violations, Inspections, and Abnormal Occurrences," *Journal of Political Economy* 97, 115–154.

Gonsalves, J. B. (1987). "Census and Analysis of the Population of Continuous, Real-Time Simulators in the Nuclear Power Plant Industry." In *Proceedings of the Society for Computer Simulation Simulators Conference*. San Diego, CA: Society for Computer Simulation.

Horowitz, J. L., and G. R. Neumann (1992). "A Generalized Moments Specification Test of the Proportional Hazards Model," *Journal of the American Statistical Association* 87, 234–240.

Jordan, E. L. et al. (1988). "Safety System Challenges in US Commercial Power Reactors." In *Nuclear Power Performance and Safety*, Volume 2: *Achievements in Construction and Operation*. Vienna: International Atomic Energy Agency.

Joskow, P., and G. A. Rozanki. (1979). "The Effects of Learning by Doing on Nuclear Plant Operating Reliability," *Review of Economics and Statistics* 61, 161–168.

Kalbfleisch J. D., and R. L. Prentice. (1980). *The Statistical Analysis of Failure Time Data*. New York: John Wiley & Sons.

Krautman, A. C., and J. L. Solow. (1992). "Nuclear Power Plant Performance: The Post Three Mile Island Era," *Energy Economics* 14 209–216.

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

Lipford, B. L., N. M. Cole, and T. J. Friderichs. (1991). "Three Mile Island Revisited," *Mechanical Engineering* 113(1), 42–47.

Marcus, A. A. (1988). "Implementing Externally Induced Innovations: A Comparison of Rule-Bound and Autonomous Approaches," *Academy of Management Journal* 31, 235–256.

Moore, C. E. (1989). "INPO Programs for Improvement of Nuclear Power Plant Operations." In *Proceedings of an NEA Symposium on Good Performance in Nuclear Projects*. Paris: OECD.

Moreau, T., J. O'Quigley, and M. Mesbah. (1985). "A Global Goodness of Fit Test for the Proportional Hazards Regression Model," *Applied Statistics* 34, 212–218.

Morimoto, K. (1986). "Comparing International Scram Frequency Statistics," *Nuclear Engineering International* (September), 43–42.

Murley, T. E. (1990). "Developments in Nuclear Safety," *Nuclear Safety* 31, 1–9.

Murphy, J. A., M. A. Cunningham, and M. T. Leonard. (1988). "NUREG-1150: The Reactor Risk Reference Document Results and Regulatory Applications." In *Nuclear Power Performance and Safety*, Volume 4: *Safety Technology*. Vienna: International Atomic Energy Agency.

Pate, Z. T. (1988). "U.S. Nuclear Utility Industry Today: Progress Since TMI," *Nuclear Engineer* 29, 171–174.

Perrow, C. (1984). *Normal Accidents: Living with High Risk Technologies*. New York: Basic Books.

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

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

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

Schoenfeld, D. (1980). "Chi-Square Goddness of Fit Tests for the Proportional Hazards Regression Model," *Biometrika* 69, 145–153.

Sturm, R. (1991). "Reliability and Maintenance in European Nuclear Power Plants: A Structural Analysis of a Controlled Stochastic Process." Unpublished doctoral dissertation, Stanford University.

Sturm, R. (1993). "Nuclear Power in Eastern Europe: Learning or Forgetting Curves?" *Energy Economics* 15, 183–188.

Thomas, S. D. (1988). *The Realities of Nuclear Power: International Economic and Regulatory Experience*. Cambridge: Cambridge University Press.

Tsiatis, A. A. (1981). "A Large Sample Study of Cox's Regression Model," *The Annals of Statistics* 9, 93–108.

U.S. Nuclear Regulatory Commission (1975). *Reactor Safety Study: An Assessment of Accident Risks in U.S. Commercial Nuclear Power Plants*, NUREG-75/014 (WASH-1400). Washington: U.S. Government Printing Office.

U.S. Regulatory Commission. (monthly from January 1976 to December 1991). *Licensed Operating Reactors Summary Status Report* (NUREG-0020). Washington: U.S. Government Printing Office.

U.S. Regulatory Commission. (1980a). *NRC Action Plan Developed as a Result of the TMI-2 Accident* (NUREG-0660). Washington: U.S. Government Printing Office.

U.S. Regulatory Commission. (1980b). *Clarification of the TMI Action Plan Requirements* (NUREG-0737). Washington: U.S. Government Printing Office.

U.S. Regulatory Commission. (1981). *Guidelines for Control Room Design Reviews* (NUREG-0700). Washington: U.S. Government Printing Office.

U.S. Regulatory Commission. (1989). *Severe Accident Risks: An Assessment for Five U.S. Nuclear Power Plants* (NUREG-1150). Washington: U.S. Government Printing Office.

Warrock, A. (1987). "Boston Edison's Pilgrim Plant: New Blood, Persistent Problems," *New England Business* 9(14), 39–40.

Wiegle, H. R. (1986). "Nuclear Plant Scram Reduction." In *Instrumentation in the Power Industry: Proceeding of the Twenty-ninth Power Instrumentation Symposium*. Research Triangle Park, NC: ISA.

Yates, M. (1988). "TVA: A Need for Accountability?" *Public Utilities Fortnightly* 121(10), 39–41.