



Managing Safety-Related Disruptions: Evidence from the U.S. Nuclear Power Industry

Christian C. Blanco,^{1,*} Felipe Caro,² and Charles J. Corbett²

Low-probability, high-impact events are difficult to manage. Firms may underinvest in risk assessments for low-probability, high-impact events because it is not easy to link the direct and indirect benefits of doing so. Scholarly research on the effectiveness of programs aimed at reducing such events faces the same challenge. In this article, we draw on comprehensive industry-wide data from the U.S. nuclear power industry to explore the impact of conducting probabilistic risk assessment (PRA) on preventing safety-related disruptions. We examine this using data from over 25,000 monthly event reports across 101 U.S. nuclear reactors from 1985 to 1998. Using Poisson fixed effects models with time trends, we find that the number of safety-related disruptions reduced between 8% and 27% per month in periods after operators submitted their PRA in response to the Nuclear Regulatory Commission's Generic Letter 88-20, which required all operators to conduct a PRA. One possible mechanism for this is that the adoption of PRA may have increased learning rates, lowering the rate of recurring events by 42%. We find that operators that completed their PRA before Generic Letter 88-20 continued to experience safety improvements during 1990–1995. This suggests that revisiting PRA or conducting it again can be beneficial. Our results suggest that even in a highly safety-conscious industry as nuclear utilities, a more formal approach to quantifying risk has its benefits.

KEY WORDS: Econometrics; nuclear energy; probabilistic risk assessment; safety

1. INTRODUCTION

The direct and indirect benefits of risk assessments are difficult to measure, especially for low-probability, high-impact events. The direct benefits of risk assessments are hard to estimate because the counterfactual is rarely observable, and in many cases, a metric for measuring safety outcomes is not obvious. The indirect benefits of risk assessments are difficult to quantify because the link between assessments and improvements is not always clear. These

are some of the reasons why managers may not invest time in conducting quantitative risk assessments.

The importance of risk management becomes salient after a major accident, prompting managers to look for ways to prevent a recurrence. For instance, the 2010 Deepwater Horizon incident, the largest oil spill in the United States, prompted the offshore oil drilling industry to search for new risk assessment tools. The National Commission on the BP Deepwater Horizon Oil Spill (2011, p. 251) recommended looking at the risk management experience of the nuclear industry. Several scholars have explored the potential role of quantitative risk assessments in the offshore oil drilling industry (Paté-Cornell, 1995; Paté-Cornell & Bea, 1992), yet the adoption of PRA in this industry remains low. The operational data on nuclear power production and the 40 years of

¹Fisher College of Business, Ohio State University, OH, USA.

²Anderson School of Management, University of California, Los Angeles, CA, USA.

*Address correspondence to Christian Blanco, Department of Management Sciences, Fisher College of Business, Ohio State University, OH 43210, USA; blanco.58@osu.edu.

experience with probabilistic risk assessment³ (or PRA) in this industry can provide insights on cost-effective risk management tools for the offshore oil drilling industry (Azizi, 2014; Cooke, Ross, & Stern, 2011). The purpose of PRA is to quantify the likelihood and consequences of accidents in operating complex technologies. The challenges of assessing risks in offshore oil drilling are comparable to those in the nuclear industry because both industries face low-probability, high-impact events in their operations (Cooke et al., 2011; National Commission on the BP Deepwater Horizon Oil Spill, 2011, p. 235; Paté-Cornell, 1995). Yet, the impact of PRA on safety performance is still not well measured even in the nuclear sector (Goerlandt, Khakzad, & Reniers, 2016).

The economic pressures at nuclear plants often dictated the management priorities (MacAvoy & Rosenthal, 2005), and the short-term goals may discourage managers from focusing on the potential long-term benefits of PRA. Early works on safety in the offshore oil industry and in the nuclear industry show how many of the underlying causes of major safety issues are rooted in management practices. In an analysis of the Piper Alpha accident, Paté-Cornell (1995) describes how three out of the four root causes of the platform fire are related to management (i.e., personnel issues, economic pressures, and maintenance practices) and only one is rooted in design (see fig. 5, p.106). The management practices (e.g., maintenance priorities, conducting PRA) at nuclear plants were not standardized in the early decades of the industry and data on past failures were hard to obtain (Carroll, Sterman, & Marcus, 1998; Rees, 1996). We explore the mechanisms of how PRA may have facilitated changes in management practices that improved overall safety at nuclear plants.

It is not clear that managers should expect any operational benefits (not directly related to safety) from implementing PRA for a couple of reasons. First, it is not obvious that managers would experience any further improvements given the already strict oversight from regulators and the existing focus on safety associated with the production of nuclear power. Second, the focus of PRA is on modeling and quantifying risks, not on improving operations. Yet, some argue that the benefits of PRA can go beyond producing a numerical assessment of risk (Apostolakis, 2004; Pence et al., 2018).

There are potential cost savings benefits from PRA. For instance, Paté-Cornell (1990) estimates that the cost associated in implementing safety improvements from PRA can be two orders of magnitude lower compared to the cost of reaching the same level of safety through structural changes. There are also potential benefits to plant capacity factors due to enhanced reliability of components and systems. (See NRC Report NUREG-1560 [1997] for a comprehensive list of maintenance-related improvements that contributed to improved reliability and productivity.) Despite many articles describing the benefits of PRA, the perception of PRA's usefulness is mixed because the probabilities obtained from these risk assessments are subjective (Bier, 1999; Goerlandt et al., 2016) and the models used to analyze them may be incomplete (Rae, Alexander, & McDermid, 2014). The mixed perception of PRA may influence how much value practitioners can get from it (Goble & Bier, 2013). The widespread variation in PRA adoption is indicative that its safety and operational benefits are not so apparent.

Examining the U.S. nuclear industry allows us to explore the impact of risk assessments on preventing safety-related disruptions. We conducted 17 background interviews with members of the industry and PRA experts to guide our data collection. Our interviews reveal the lack of empirical analysis of the impact of individual plant assessments on preventing unusual, safety-related events. Our article is closely related to Pence et al. (2018), who measure the monetary value of PRA with industry data. Our focus is on the early developments of PRA when it may have had the largest impact. There are studies that demonstrate how PRA can identify ways to improve safety and operations at the plant, but most of these papers quantify improvements based on the numbers produced by the PRA, not on actual performance. The purpose of this article is to explore the impact of PRA adoption on the frequency of unusual events, not the numerical estimates of PRA itself.

The contribution of this article is to explore the impact of PRA on safety using observable outcomes. (This is in the spirit of performance-based regulation, that is, assessing risks based on observable results, rather than estimating the risks of potential disasters that may not be observed. See the Nuclear Regulatory Commission [NRC, 2018] for a description of performance-based regulation.) Although the metrics we examine here are different than the performance indicators used by the NRC in

³See <https://www.nrc.gov/about-nrc/regulatory/risk-informed/pr.html> for a current definition of PRA; last visited on November 5, 2018.

its implementation of performance-based regulation, our study identifies the impact of PRA on observable outcomes. We explore this using data from over 25,000 safety-related event reports at nuclear power plants for periods before and after operators submitted their risk assessments to the regulators. We focus on the period from 1985 to 1998 because this is when PRA was widely adopted by the nuclear industry; it has been updated several times since then. Our estimate of the effect of PRA on the nuclear power sector does not lead to predictions for other areas. However, given the preexisting focus on safety in nuclear power, it is plausible that the effect of conducting PRA in other sectors would be at least as large as what we find in this industry.

Our results show that operators experienced between 0.16 and 0.54 (or 8%–27%) fewer safety-related events per month compared to an average of two per month in periods before they submitted their risk assessments. To put this number in perspective, the occurrence of one safety-related event is associated with, on average, a 7.4% decrease in production capacity for that month.⁴ This impact is significant even after controlling for cumulative experience, voluntary shutdowns, and regulatory penalties (forced shutdowns). PRA can facilitate better information flow of operating experience and faster learning of best practices (Paté-Cornell, 1990, pp. 1211 and 1215); we find evidence that even a one-time adoption of PRA can stimulate learning and decrease the rate of recurring events. We examined the abstracts of the safety reports and identified whether they reference a similar event, suggesting that issue is recurring. We find that recurring events decreased from an average rate of 4.3% in periods before conducting PRA to an average of 1.9% after. This suggests that PRA was effective in promoting information flow and sharing of operational experience within the industry.

We conduct several robustness tests to explore the variation in the impact of PRA across plants and on different types of events. The magnitude of the impact of PRA varies across different plants and by the type of significant (safety) event. The estimates we obtain are approximate and are far from universal, but the results across all our main results and robustness tests show a negative and statistically significant relationship between PRA submission and safety-related disruptions.

⁴See the Appendix for the association between safety-related events and electricity generation.

The outline of the article is as follows. In Section 2, we briefly describe the history of PRA within the U.S. nuclear power industry and discuss its broader benefits beyond providing a numerical value of risk. We present the data in Section 3, followed by the methods and results in Section 4. We present robustness tests in Section 5. Our discussions and conclusions are in Section 6.

2. PRA IN THE NUCLEAR INDUSTRY

What is the impact of PRA on the frequency of safety-related events? The time line of how PRA developed in the nuclear power industry reflects some of the challenges in establishing its benefits beyond quantifying a probability of a meltdown. Several studies suggest that conducting PRA had benefits that extended to overall operations. We will revisit some of those studies here, showing the lack of an industry-wide measure on the impact of PRA on operational performance. At the end of this section, we explain why we limit our study to 1985–1998. A technical description of PRA is available in the NRC Report NUREG/CR-6823 (2003) and a more elaborate historical background of PRA can be found in Keller and Modarres (2005).

2.1. Historical Perspective of PRA in U.S. Nuclear Power

The slow diffusion of PRA within the nuclear power industry reflects some of the uncertainty about its benefits. The earliest application of PRA at U.S. nuclear plants was published in October 1975 (NRC Report WASH-1400). In 1979, the Three Mile Island Reactor 2 had a core meltdown due to loss of cooling water. This incident led the industry and the regulators to explore more widespread use of better risk management tools. In 1983, the regulators published the first PRA procedures guide (NRC Report NUREG/CR-2300). In 1988, the regulators released Generic Letter 88-20 requiring all nuclear reactors to conduct “Individual Plant Examinations for Severe Accidents Vulnerabilities.” All plant owners decided to fulfill the requirements of the individual plant examinations using PRA⁵ (NRC Reports NUREG-1150, 1991, p. 1, and NUREG/

⁵We are grateful to Professor Mohammed Modarres at the University of Maryland for pointing us to Appendix A in NUREG 1560, the document where we obtained individual PRA adoption dates.

CR-5750, 1999; Lochbaum, 2000, p. 4). We focus on this industry-wide adoption of PRA to estimate the impact of PRA on safety performance.

There are operators that conducted PRA before the NRC requested it in Generic Letter 88-20, so their PRA submission to the NRC is not their first opportunity to conduct PRA. We conduct tests with the set of operators that conducted PRA for the first time in response to Generic Letter 88-20 and those that completed PRA before the NRC required it. One of our tests will examine the impact of doing PRA for the first time, and we will also test whether revisiting PRA or conducting it again has any benefits at all. We also estimate the impact of PRA with a larger, industry-wide sample.

In the early 1980s, there was significant variation in the safety and operations at nuclear power plants (NRC Report NUREG/CR-5750, 1999; Rees, 1996, pp. 100–103), but the industry was improving over time. The variation in safety and operating performance was so dramatic that this information was used at industry-wide meetings with executives of nuclear utilities to shame the worst performers (Rees, 1996, pp. 103–105). Some of these industry-wide initiatives to improve safety include the formation of the Institute of Nuclear Power Operators (INPOs) (Taylor & Wolak, 2012), changes within the states' public utilities commissions (Fremeth & Holburn, 2012), and shifts in risk perception. Wheatley, Sovacool, and Sornette (2017) do observe the overall decline in significant events globally, but they do not directly measure what portion of this decline is explained by PRA. These industry-wide trends in safety over a long horizon create challenges for estimating the impact of PRA, but these risk assessments are a specific intervention during a well-defined window, making it possible to identify the effect of PRA.

2.2. Case Studies on the Operational Benefits of PRA

Early practitioners of PRA claim that the numeric aspect of PRA is not the end itself, but that the process of doing PRA leads to other operational benefits (Garrick, 1989). For instance, Dubord, Golay, and Rasmussen (1996) describe how they use PRA to prioritize maintenance and inspection resources at the Fitzpatrick plant. PRA was also useful in developing in-service inspections at the Surry 1 plant as described in Vo et al. (1993). Better allocation of inspection resources and conducting maintenance during operations can lead to fewer safety-related issues

and higher productivity (Carroll et al., 1998). The NRC describes a list of common plant improvements from PRA in Table 2.3 in NUREG-1560. (pp. 2–11 and 2–12), 45% of which are low-cost, operational changes as opposed to changes in design and hardware. There are also examples in the off-shore oil industry that explore the benefits of PRA in facilitating learning and sharing of operational experience (Pate-Cornell, 1990).

We draw on the experience of two more plants to show how PRA may have accelerated learning, reducing recurring events at the plant. Sharing their experience of using PRA at the Arkansas 1 plant, Vo, Harris, and Gore (1989) summarize the different conditions that could lead to equipment failure (Table II.A) and how to manage them (Table II.B). This type of knowledge can allow operators to avoid these conditions and thus reduce the number of safety-related events at the plant. In a joint study with the Electric Power Research Institute, Worledge and Wall (1989) describe how they use PRA for the Seabrook plant and past event records across many plants to create a common-cause failure database. This is an example of how performing PRA can be an opportunity to examine past experience and to learn from it, reducing the recurrence of safety issues.

Many of these studies describe the changes associated with PRA but did not quantify the benefits, albeit with some exceptions. The work by Daling, Marler, Vo, Phan, and Friley (1995) is one of the few that quantify the benefits to changes in heating, cooling, and water systems from PRA. These studies focus on the PRA models and the changes implemented at those plants, but they do not exploit industry-wide variations to systematically compare performance outcomes in periods before and after conducting the PRA.

2.3. A Focus on the Time of Submission as Proxy for Adoption of PRA

We limit our study to the period 1985–1998 for several reasons. First, this is the first time PRA was required in the U.S. nuclear sector; it has been updated several times since, but including later periods would make it harder to identify the effect of initial adoption. Second, the industry-wide implementation of PRA occurred during 1990–1995. Many operators conducted PRA before 1990, but we include that subset in most of our tests to explore whether the process of reviewing or repeating PRA has a substantial impact on safety. Third, the reporting

guidelines for unusual events were standardized after 1984 (NRC Report NUREG-1022); this minimizes differences in variation due to reporting. Fourth, we do not include events in periods after 1999 because the nuclear power industry was experiencing significant market and regulatory changes; some nuclear reactors were privatized but some remained public (see Davis & Wolfram, 2012a). Fifth, this time period gives us enough variation to compare periods before and after nuclear operators adopted PRA. Our interviews with current and former members of the industry, PRA experts, and nuclear regulators reveal contemporaneous industry-wide changes that may also contribute to improvements in safety over time, such as those we mentioned in Section 2.1. These changes contribute to improvements in safety and productivity over time, making it difficult to test the impact of PRA on safety. We discuss our tests to overcome some of these challenges later.

We draw an analogy to other studies that measure the impact of a management tool on firm performance. Using the adoption of ISO 9000⁶ from 1987 to 1997, Corbett, Montes-Sancho, and Kirsch (2005) estimate that a firm’s initial ISO 9000 certification is associated with significant financial benefits. Examining quality management awards given to 463 firms from 1983 to 1993, Hendricks and Singhal (1997) measure the impact of having a successful quality management system on sales growth. There was a lot of skepticism around the effectiveness of these management tools during their early stages because, just like PRA, the benefits of adopting these programs are not easy to quantify.

3. DATA ON U.S. NUCLEAR EVENT REPORTS AND ADOPTION OF RISK ASSESSMENT FROM 1985 TO 1998

We collected all event reports (formally referred to as Licensee Event Reports or LERs) submitted by nuclear operators from 1985 to 1998. A reportable event is determined by Title 10, Code of Federal Regulations (10 CFR). These reports describe an event with significant safety implications at the plant. All event reports are publicly available on the nuclear regulator’s website.⁷ Based on our interview with a member of the industry, the primary purpose of an

Table I. Summary Statistics of Reactor and Monthly Data

Statistic	Mean	St. Dev.	Min	Max
Monthly events	1.614	1.731	0	16
Before risk assessment	1.961	1.903	0	16
After risk assessment	1.204	1.395	0	13
Rate of recurring events	0.03	0.15	0	1
Before risk assessment	0.04	0.17	0	1
After risk assessment	0.02	0.12	0	1
Days between events	18.452	24.244	0	200
Before risk assessment	15.369	20.069	0	200
After risk assessment	24.401	29.852	0	197
Monthly Observation	Mean	St. Dev.	Min	Max
Refueling	0.079	0.270	0	1
Long-term shutdown	0.042	0.199	0	1
NRC-forced shutdown	0.016	0.127	0	1

Note: Operators experience 1.6 events per month, of which 3% were recurring. The average number of monthly events before PRA is higher at 2 per month compared to the 1.2 average after PRA. The rate of recurring events before PRA is also higher at 4% compared to the overall rate. The relative decrease of the rate of recurring events before and after PRA is 50%; we will measure this more precisely using Poisson regression models in later sections.

event report is to record operating experience at a nuclear plant to serve as a learning tool for nuclear power operators and the regulators. Event reports are used to collect data on the failure rates of various components. They are also recorded to document various events that activate safety equipment or that may prevent safety equipment from functioning properly. There are different types of reportable events, but we initially treat them as equivalent. We look at different types of events separately later. Although event reports may be used to aid the regulator’s oversight and compliance rules, our interview with a nuclear power regulator confirms that they are not the basis for levying fines or penalties against reactors. The events analyzed in PRA include those described in the reports, such as a reactor trip or equipment failure, but PRA is not limited to those.

Operators submitted their risk assessment reports in different years (Appendix A of the NRC Report NUREG-1560). Regulators required plants to conduct their own risk assessments with the intent of increasing their appreciation for risk management (NRC Generic Letter 88-20, 1988).

Table I shows the summary statistics. The average monthly number of events is higher in periods before operators submitted their PRA to the NRC, at 1.9 per month, compared to periods after submission, at 1.2 per month. The overall rate of recurring

⁶ISO 9000 is an international quality management systems standards.

⁷<https://lersearch.inl.gov/LERSearchCriteria.aspx>, last accessed March 10, 2016.

events⁸ is very low at 3%, but we see that the rate is 2% in periods after submission compared to 4% in periods before that. A related metric, the average number of days between two successive events, is 18.5 days. This metric captures reliability because it is the time between two failures. We focus on the number of events per month, but we also explore the impact of PRA on reliability in Section 5.5.

We included monthly data on operational and regulatory levers that may influence the average monthly frequency of disruptions at the plant. Reactors refuel every 12 to 18 months, and this takes anywhere from 30 to 60 days. David Lochbaum, the director of the Nuclear Safety Project, provided us with the data on when reactors refuel. Table I shows that reactors spend 7.9% of all months refueling. Based on our interviews, reactors perform maintenance and process improvements during refueling, so the cumulative number of times reactors have refueled may influence the frequency of monthly events. We also collected data on long-term voluntary plant shutdowns, which can take more than a year.⁹ Operators implement improvements at the plant at these times, so we try to control for them in our regression models.

Regulators imposed a forced shutdown in only 1.6% of months in the sample. Based on our interviews, these penalties are associated with significant changes in both management and process improvements at the plant, so we also control for this by creating a variable that captures the cumulative number of forced shutdowns over time.¹⁰

We have additional data on reactors, such as capacity, manufacturer, and the regulatory region. We include fixed effects by plant, so controls that do not vary over time are not included here.

4. METHODS AND RESULTS

We want to test the impact of PRA on the number of monthly safety-related events. In this section, we describe the regression model we use, then present the results on the impact of PRA on monthly events, the rate of change, and a test that explains one of several possible causal mechanisms.

⁸We describe how we identify recurring events in Section 4.3.

⁹There is a high opportunity cost with shutting down the plant. To put this in perspective, a 1,000 MW plant that operates at an average price of \$80 per MWh can lose up to \$60 million in revenue a month.

¹⁰We are grateful to Adam R. Fremeth at the Ivey Business School at Western University for providing us historical records of forced shutdowns.

4.1. What Is the Impact of PRA on Safety-Related Events?

We estimate the impact of PRA on monthly events using a Poisson regression equation. We model the (random) number of monthly events for reactor i at time t as Y_{it} using a Poisson distribution because the number of monthly events takes discrete values from 0, 1, 2, ... and so on. We denote the expected number of events as λ_{it} . We assume that the expected number of events has a log-linear relationship with a set of explanatory variables \mathbf{X}'_{it} such that $\ln\lambda_{it} = \mathbf{X}'_{it}\boldsymbol{\beta}$, where $\boldsymbol{\beta}$ is a set of coefficients that will be estimated. See Greene (2012) for more details on this formulation. The number of (observed) events y_{it} that occur for reactor i at time t conditional on reactor and time-specific variables \mathbf{X}_{it} is then:

$$P(Y_{it} = y_{it} | \mathbf{X}_{it}) = \frac{e^{-\lambda_{it}} \lambda_{it}^{y_{it}}}{y_{it}!}, \quad y_{it} = 0, 1, 2, \dots \quad (1)$$

Therefore, the expected number of events y_{it} is given by $E[y_{it} | \mathbf{X}_{it}] = \lambda_{it} = e^{\mathbf{X}_{it}\boldsymbol{\beta}}$. We estimate this equation using R software and the package *poissonmfx* for marginal effects.

We want to see whether the adoption of PRA is associated with changes in the average number of safety-related events. We create a binary variable *Submitted PRA* _{it} equal to 1 on month t and subsequent months when reactor i submitted its PRA and 0 otherwise. Because we are interested in estimating changes within a reactor, we include a dummy variable δ_i for each reactor i . We also include a linear year trend $year_t$ to control for industry-wide trends. We can include other time-varying variables, denoted by $\mathbf{X}'\eta$. We estimate the following Poisson regression equation:

$$\begin{aligned} \text{Mean Events}_{it} = & \exp(\beta_1 \times \text{Submitted PRA}_{it} \\ & + \beta_2 \times year_t + \mathbf{X}'\eta + \delta_i). \end{aligned} \quad (2)$$

The Poisson regression results, displayed as marginal effects, of Equation (2) are summarized in Table II. This sample includes 101 nuclear reactors. The estimate of PRA is associated with a 0.23 decrease ($p < 0.05$) in monthly events after including the cumulative number of times operators refueled the reactor. The estimate of PRA is associated with a decrease of 0.22 and 0.23 even after controlling for voluntary and regulator-forced shutdown. The results of Table II confirm that the adoption of PRA is associated with a decrease in monthly events. These

Table II. Poisson Regression of PRA Adoption on Average Monthly Events Across 101 Nuclear Reactors

	Dependent Variable: Monthly Events		
	(1)	(2)	(3)
Submitted PRA	-0.234** (0.093)	-0.219** (0.093)	-0.231** (0.093)
Reactor fixed effects	Yes	Yes	Yes
Time trends	Yes	Yes	Yes
Reactor-level trends	Yes	Yes	Yes
Cumulative refueling months	Yes	Yes	Yes
Cumulative long-term shutdown	–	Yes	Yes
Cumulative NRC-forced shutdown	–	–	Yes
Observations	16,066	16,066	16,066
McFadden's R^2	0.10	0.10	0.10

Note: Cluster-robust standard errors are in parentheses. We report the marginal effects of the results. It is possible that the estimates we have provided are still biased due to potential omitted variables or self-selection issues. The standard errors (and McFadden's R^2) are the same when rounded but differ at the third decimal place.
 * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

estimates provide the average impact of PRA across the industry.

Now, we examine the impact of PRA between operators that conducted their first PRA in response to Generic Letter 88-20 and those that implemented PRA before it was required by the NRC. (The list of operators that completed a PRA before the introduction of Generic Letter 88-20 is in the NRC Report NUREG-1050.) Models (1)–(3) in Table III show the results for the sample of operators that conducted their first PRA in response to Generic Letter 88-20. We refer to these operators as the reactive adopters. Models (4)–(6) show the results for the sample of proactive adopters, operators that completed PRA before Generic Letter 88-20.

Model (1) in Table III shows the Poisson regression results with reactor fixed effects, linear time trends, and controls for the cumulative number of times operators have refueled. We find that the number of monthly events decreased by 0.17 ($p < 0.01$) in periods after the adoption of risk assessment or roughly 8% relative to the average before adoption. The impact of PRA in Models (2) and (3) remains statistically significant and its magnitude is similar to Model (1) even after adding controls for the cumulative number of voluntary and forced long-term shutdowns.

We compare the impact of PRA with the subsample of operators that conducted PRA before Generic Letter 88-20 to those that adopted it for the first time when it was required by the NRC. The results for that subsample are summarized in Models (4)–(6) in Table III. Model (4) shows the results with the same controls as in Model (1). We see that the impact of PRA is associated with a 0.53 ($p < 0.01$) decrease in average monthly number of events or roughly a 27% decrease relative to the overall average number of events before any of the operators submitted their PRA to the NRC. The magnitude of the impact is about the same if we add controls for the cumulative number of voluntary shutdowns in Model (5), if we include controls for the cumulative number of forced long-term shutdowns in Model (6). The results show that the impact of submitting a PRA to the NRC is about three times larger for those that conducted PRA previously. Overall, we find evidence that the impact of these quantitative risk assessments are associated with roughly an 8%–27% decrease in the average number of monthly events. We provide a short discussion of this result in Section 6.1.

4.2. Does the Rate of Decrease of Monthly Events Change After Submitting PRA?

Given that there is an industry-wide decreasing trend in the average number of events, we need to test whether the rate of decrease in the number of events changes after the submission of PRA. To capture this, we create the variable “Monthly trend before submission” as follows. Suppose a reactor started operating in January 1985 and filed its PRA in January 1990. The monthly trend before filing takes value 1 in January 1985, 2 in February 1985, and so on, and it takes 0 when it files its PRA and after. We create the variable “Monthly trend after submission” in a similar fashion: it takes value 1 in January 1990, 2 in February 1991 and so on, and 0 in all months before it submitted its risk assessment. We work with the industry-wide sample henceforth to examine the impact of PRA on the rate of decrease of monthly events and on preventing recurring events.

Table IV shows that the monthly rate of decrease is faster in periods after operators submitted their risk assessments. Model (1) shows the estimates after including year dummy variables. PRA adoption is associated with a 0.75 decrease in average monthly events. The result shows the rate at which operators reduced the number of events by 0.02 ($p < 0.01$) per

Table III. Marginal Effects of PRA Adoption on Average Monthly Events Across Reactive and Proactive Adopters

	Dependent Variable: Monthly Events					
	Sample of Reactive Adopters ^a			Sample of Proactive Adopters ^b		
	(1)	(2)	(3)	(4)	(5)	(6)
Submitted PRA	-0.165*** (0.047)	-0.149*** (0.047)	-0.158*** (0.047)	-0.529*** (0.088)	-0.521*** (0.088)	-0.538*** (0.092)
Reactor fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Time trends	Yes	Yes	Yes	Yes	Yes	Yes
Reactor-level trends	Yes	Yes	Yes	Yes	Yes	Yes
[†] Refueling months	Yes	Yes	Yes	Yes	Yes	Yes
[†] Long-term shutdown	–	Yes	Yes	–	Yes	Yes
[†] NRC-forced shutdown	–	–	Yes	–	–	Yes
Observations	13,204	13,204	13,204	2,862	2,862	2,862

Note: Cluster-robust standard errors are in parentheses. We report the marginal effects of the results. We use NUREG-1050 to identify the sample of operators that conducted PRA before the NRC requested it in Generic Letter 88-20. See Tables 3-1, 3-2 on pages 3-7, 3-8 in NRC Report NUREG-1050 for that list.

^aThis is the subsample of operators that adopted PRA for the first time in response to Generic Letter 88-20.

^bThis is the sample of operators that completed PRA before it was required by the NRC.

[†]These variables are cumulative.

* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

Table IV. Rate of Decrease of Events per Month, Before and After Filing PRA, Using Poisson Regression

	Dependent Variable: Monthly Events
	(1)
Submitted PRA	-0.752*** (0.274)
Monthly trend before submission	-0.009*** (0.003)
Monthly trend after submission	-0.021*** (0.008)
Reactor fixed effects	Yes
Year (dummies)	Yes
Observations	16,066

Note: The marginal effects of the results are provided. Cluster-robust standard errors are provided in parentheses.

* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

month in periods after submitting their PRA, compared to 0.01 ($p < 0.01$) per month before submission. The results suggest that PRA not only reduces the number of events, but also accelerates the rate at which plants improve their safety performance.

The result in Table IV suggests that PRA is effective in reducing the average frequency of events as well as accelerating the rate of improvement over time. Fig. 1 is a graphical representation how our models measure the average decrease in the number of events and the faster rate of improvement in

the presence of a downward-sloping trend. Fig. 1(A) depicts the average decrease after PRA submission with a downward sloping trend. Fig. 1(B) illustrates the average decrease in the number of events as well as a steeper rate of decrease after PRA submission.

4.3. The Impact of PRA on Recurring Events

In the previous subsections, we show evidence that PRA is associated with a decrease in safety-related events, but we have not yet explored causal mechanisms that may explain why. It is possible that several mechanisms exist; the purpose of this subsection is to explore one of them. Operators often revisit past event reports when they conduct PRA (Worledge & Wall, 1989). Conducting PRA may give operators an opportunity to gain more knowledge from LERs by identifying recurring, common-cause failures and to explore ways to prevent them.

Operators are required to report similar events in the abstract of an LER. The abstracts mention the term “LER(s),” followed by a numerical code to reference a past event. We use text analysis on the abstracts of the LERs to check if a similar event occurred in the past. We removed all punctuation in the LERs (e.g., converting LER’s to LERs), then we use text processing to code an event as recurring if the term “LER” is present, followed by a sequence of numbers.

Fig. 1. Average versus the rate of decrease in the number of events before and after PRA.

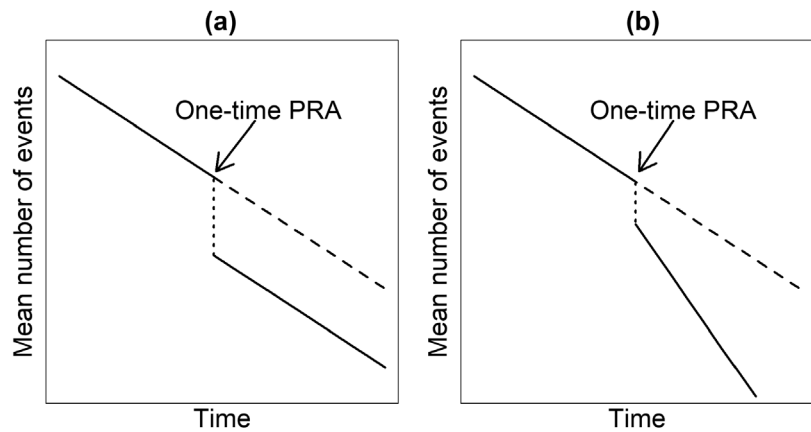


Table V. The Impact of PRA on the Rate of Recurring Significant Events

Dependent Variable: Ratio of Recurring Events to Total Events			
	(1)	(2)	(3)
Submitted PRA	-0.024*** (0.003)	-0.025*** (0.003)	-0.025*** (0.003)
Reactor fixed effects (δ_i)	-	Yes	Yes
Adoption is lagged for one year	-	-	Yes
Observations	11,262	11,262	11,262
Adjusted R^2	0.006	0.054	0.055

* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

We estimate the impact of PRA on the ratio of recurring events to the total number of events for that month. The average rate of events that recur is 4.3% in months before any PRA was submitted to the regulators. (Months without a significant event are not included in this analysis, and therefore our models do not include time trends.) Table V shows that PRA is associated with a decrease of 2.5 percentage points from that average. This change translates to a 42% decrease in the rate of recurring events in periods with PRA.

5. ROBUSTNESS TESTS

So far, we have presented evidence of a significant association between adopting PRA and a reduction in the frequency of events as well as an acceleration in the rate of reduction of significant safety events. We have not examined any possible indirect costs or benefits of doing PRA on capital additions and production costs. There are also potential biases that we have not yet addressed.

In this section, we explore the impact of PRA on capital and production costs. Then, we explore the variation of the impact of PRA on the best and worst performing plants and on different types of events. This is followed by two robustness tests using an alternative metric.

We do these extensions for several reasons. First, it is not clear whether PRA is associated with an increase in capital or production costs. Second, the benefits of PRA may vary across plants and different types of events. Third, very few industry-wide studies exist that explore the impact of PRA on reliability, measured in time between two safety-related events.

5.1. The Impact of PRA on Capital Additions and Production Costs

We use annual capital additions and production costs from a report published by the Oak Ridge National Laboratory (NRC Report NUREG/CR-6577, 2003). The data show that the average production costs increased from 1980 to 1990, but that the average stops increasing around the same time most nuclear power operators started implementing PRA in response to Generic Letter 88-20. One possible explanation for the increase in cost from 1980 to 1990 is the industry’s response to the Three Mile Island incident in 1979. Managers implemented many changes and retrofits in response to new regulatory requirements; many of these efforts were not voluntary. There are several other factors that may contribute to production costs increasing or decreasing around 1990, and those include policies and institutions introduced around that time such as the maintenance rule, the INPO, and the industry-wide adoption of PRA. This is why we include industry-wide trends as

Table VI. The Impact of PRA on Capital and Production Costs
(in Millions of Dollars in 2001 Values)

	Dependent Variable	
	Capital Costs (1)	Production Costs (2)
Submitted PRA	-7.304 (12.889)	-20.576*** (4.113)
Reactor fixed effects	Yes	Yes
Time trends	Yes	Yes
Reactor-level trends	Yes	Yes
Cumulative refueling months	Yes	Yes
Cumulative long-term shutdown	Yes	Yes
Cumulative NRC shutdown	Yes	Yes
Observations	1,319	1,319
R ²	0.232	0.889

Note: The number of observations here is lower because we only have annual data on capital and production costs. We use ordinary least squares to measure the association of PRA on capital and production costs.

* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

a control in all of our tests. Some members of the industry were looking for alternative, cost-effective measures, and PRA was identified as a possible option. Our next test explores the association of PRA submissions in response to Generic Letter 88-20 on annual capital additions and production costs from 1985 to 1998.

We tested the impact of PRA on annual capital additions and production costs with the same set of controls in our earlier models. We use ordinary least squares (OLS) for this robustness test because the dependent variables have continuous values. The dependent variables in Models (1) and (2) in Table VI are in millions of dollars adjusted to 2001 values. Model (1) in Table VI shows no evidence that PRA is associated with an increase in annual capital additions; it is in fact negative but not statistically significant at the 0.10 level. Model (2) shows that PRA is associated with roughly a \$20.6 million decrease in annual production costs (in 2001 values).

We estimate the cost of doing PRA in 1985 and compare it to the indirect benefits we calculated. It is important to note that the cost of PRA may vary widely across plants and over time. We also caution against extrapolating these estimates to the cost of doing PRA today. The sentiment of various experts on whether the cost of PRA is higher or lower today is mixed. One reason for this is that the cost of conducting PRA is also difficult to measure because there could be hidden costs associated

with it. Examples of hidden costs can include the time it takes to develop and refine the PRA method or the lost time in productivity for training. Therefore, the comparison we are about to make is only a rough estimate of the cost effectiveness of doing PRA during its first decade of implementation. We obtain cost estimates of PRA from the U.S. General Accounting Office (1985) for a couple of reasons. First, reports submitted to the GAO are likely to be reliable, accurate estimates because the role of GAO is to ensure accountability to the public. Second, GAO sets high standards for fact-based reports, and is thus more likely to have verifiable cost accounting standards. The cost of PRA ranges from \$329,000 to \$988,000 (in 2001 values). The cost reduction of \$20.6 million per year (in 2001 values) is significantly larger than the cost of doing PRA. Although these costs and benefits may no longer be representative of PRA today, our estimates suggests that there are significant benefits in doing quantitative risk assessments even in a highly regulated industry.

We end this subsection with a few caveats. The estimates of the benefits of PRA we present here may be larger than the benefits operators experience today. One reason is that most of the benefits from PRA may have been realized in its early stages (in the 1980s and 1990s). PRA continues to evolve today, and operators are likely to reap the benefits that come with these changes. Moreover, we found that plants with prior experience with PRA saw greater reductions in safety-related disruptions than plants that first adopted in response to Generic Letter 88-20. Although the focus of our article is not on the accuracy of the economic impact of PRA, there is strong evidence that PRA was cost effective at that time. The results suggest that the safety and production benefits of PRA are at least an order of magnitude larger than the cost of implementing it.

5.2. The Impact of PRA on the Best and Worst Performing Plants

The adoption of PRA may have little to no impact on plants that already had a strong history of safety performance because they have little room for improvement. If this is the case, then our earlier approach may underestimate the impact of PRA on less well-performing plants. We test Equation (2) by quartile of the average number of events recorded in 1988, capturing individual operators' past safety performance. Operators in the first quartile, those

Table VII. Poisson Regression Results (Marginal Effects) of PRA by Quartile of Performance in 1988

	Dependent Variable: Monthly Events			
	Quartile 1 ^a (1)	Quartile 2 (2)	Quartile 3 (3)	Quartile 4 (4)
Submitted PRA	0.030 (0.145)	-0.232* (0.136)	-0.492** (0.249)	-0.371* (0.207)
Reactor fixed effects	Yes	Yes	Yes	Yes
Time trends	Yes	Yes	Yes	Yes
Reactor-level trends	Yes	Yes	Yes	Yes
Observations	3,988	4,382	3,236	3,919
Mean number of events before adoption	1.23	1.93	2.28	2.54
Mean number of events after adoption	1.15	1.19	1.26	1.24
Percent reduction	NA	12%	21%	15%
Number of reactors	24	27	20	25

^aReactors with the lowest average number of monthly events in 1988, two years before any reactors filed their risk assessments to the regulatory body. The total number of reactors does not add up to 101 because some plants were not yet in commercial operation in 1988. * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

with the lowest number of events in 1988, are the top performers, and those in Quartile 4 the worst. The results in Table VII show that the impact of PRA for those in Quartiles 3 and 4 is close to or even larger than in our earlier estimates. This suggests that PRA can be effective in improving the performance of operators with poor safety records.

5.3. What Is the Impact of PRA on Different Types of Events?

So far, we have treated the various types of “events” as equivalent. However, it is possible that the impact of PRA may vary depending on the type of event. The results in Table VIII focus on five different types of events that have the highest safety significance at the plant. These types of events are caused by either internal or external factors. Our interview with a nuclear regulator confirms that reporting of these events is standardized by the nuclear regulators and thus comparable across operators. The impact of PRA is consistent for four of these five types. The first type of event is “system actuation.” These events involve activation of an engineering safety feature.¹¹ According to Model (1), the adoption of PRA is associated with a 0.14 decrease in average monthly events. This type of event occurs on average 0.80 times per month, so the adoption of PRA is associated with an 18%

decrease in the number of system actuation events per month.

The second type is “technical specifications.” Operators are required to report events when they shut down due to operating under prohibited conditions. For instance, if a battery charger was declared to be out of service and it could not be repaired within a certain time window, then operators must shut down the plant. The adoption of PRA is associated with a 0.12 (or 16%) decrease in the number of monthly “technical specification” events. The third type is “degradation,” events associated with the deterioration of plant equipment such as pipes and safety barriers. The adoption of risk assessment is associated with a 27% decrease in this type of event. The fourth type is when components or systems fail or are “inoperable.” Risk assessment is associated with a 10% decrease in this type of event. The fifth type is associated with events that could have prevented fulfillment of a safety function. This is the only type out of the five we examine for which we find no evidence of a significant decrease after PRA adoption. One possible reason for this is that these events are more noticeable; therefore, it may be easier to prevent their recurrence without an in-depth analysis such as PRA.

5.4. What Is the Impact of PRA on Days Between Events?

So far, we used the number of monthly events to capture safety performance at the plant, but we can also examine the effect of PRA by looking at

¹¹ An example of this is when metal rods are deployed to stop any further reaction inside the core.

Table VIII. Poisson Regression Results of PRA Submission on Monthly Events by Type

	Dependent Variable: Monthly Events by Type				
	System Actuation (1)	Technical Specification (2)	Degradation (3)	Inoperable (4)	Prevent Safety Equip. (5)
Submitted PRA	-0.144*** (0.033)	-0.123** (0.048)	-0.041*** (0.014)	-0.006*** (0.002)	-0.026 (0.018)
Reactor fixed effects	Yes	Yes	Yes	Yes	Yes
Time trends	Yes	Yes	Yes	Yes	Yes
Reactor-level trends	Yes	Yes	Yes	Yes	Yes
Refuel	Yes	Yes	Yes	Yes	Yes
Long-term shutdown	Yes	Yes	Yes	Yes	Yes
NRC-forced shutdown	Yes	Yes	Yes	Yes	Yes
Observations	16,066	16,066	16,066	16,066	16,066
McFadden's R^2	0.17	0.08	0.15	0.12	0.11
Mean events before PRA	0.80	0.77	0.15	0.06	0.18
Mean events after PRA	0.24	0.57	0.20	0.03	0.12
Percent reduction	18%	16%	27%	10%	14%

* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

the mean time between failures. This metric captures the increase in reliability given that nuclear power plants are base load generators, that is, they produce electricity continuously.¹² Our analysis in the online companion Appendix shows that the occurrence of safety-related events is associated with a 7.4% decrease in electricity production for that month (sometimes referred to as capacity factor). This means that plants are more productive if the time between successive events is longer.

We estimate the impact of adopting PRA on the number of days between two successive events k and $k - 1$ for reactor i in time t , $Days_{k,k-1,i,t}$. The unit of observation is now an individual event. We estimate the following regression:

$$Days_{k,k-1,i,t} = \alpha + \beta_1 \times Submitted\ PRA_{it} + \beta_2 \times year_t + \delta_i + \mathbf{X}'\eta + \epsilon_{it}. \quad (3)$$

The various models in Table IX all provide statistical evidence of increased reliability from the adoption of PRA. Model (5) shows that the average number of days between two events increased by roughly 3.87 days (or 25%) in periods after plants submitted their PRA relative to a baseline of 15.4 days between events prior to PRA.

¹²See <https://www.eia.gov/todayinenergy/detail.php?id=30972> for a short description of nuclear power production; last visited May 1, 2017.

6. DISCUSSION, LIMITATIONS, AND FUTURE WORK

We wanted to measure the impact of PRA on the frequency of safety-related events at nuclear power plants. Although PRA has been around for more than 40 years, very few studies exist on the size of its impact on improving safety and operations. Our interviews with members of the industry and the NRC guided us in collecting data. Our results show that the adoption of PRA is associated with a 15% decrease in the frequency of monthly events. We estimate that this reduction is equivalent to the industry avoiding \$1.6 billion per year lost due to production disruptions.

6.1. Practical Insights

PRA is effective even within an industry as tightly regulated and safety conscious as the nuclear sector. We find that the impact of PRA is significant even after controlling for experience, voluntary shutdowns, and penalties (NRC-forced shutdowns). Industry-wide and reactor-specific trends over time create several challenges in estimating the impact of PRA. We included time trends, as well as tested models with reactor-specific trends, to control for potential confounding factors. Although we cannot fully control for all industry-wide improvements, our results and robustness tests remain consistent even after controlling for different time trends. Our study may not predict the impact of PRA in other industries but, if anything, we speculate that its impact may

Table IX. OLS Regression Results of the Impact of PRA on the Number of Days Between Successive Events

	Dependent Variable: Days Between Successive Disruptions				
	(1)	(2)	(3)	(4)	(5)
Submitted PRA	4.313*** (0.617)	4.248*** (0.637)	3.912*** (0.633)	3.812*** (0.635)	3.866*** (0.641)
Reactor fixed effects	Yes	Yes	Yes	Yes	Yes
Time trends	Yes	Yes	Yes	Yes	Yes
Reactor-level trends	Yes	Yes	Yes	Yes	Yes
Cumulative refueling months	–	Yes	Yes	Yes	Yes
Cumulative long-term shutdown	–	–	–	Yes	Yes
Cumulative NRC shutdown	–	–	–	–	Yes
Observations	25,872	25,872	25,872	25,872	25,872
R ²	0.086	0.125	0.126	0.126	0.126

Note: We removed observations where the number of days between two successive events is greater than 200 as we considered those potential outliers (0.01% of the data). An outlier can occur if an event was not included in the data for whatever reason, making the number of days between two events greater than it should be. This is more likely to make our results more conservative because 15 of these instances occurred in periods before PRA adoption and 40 in periods after.

* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

be larger for other industries where safety is not as highly regulated as with nuclear power production.

The 15-year gap between when PRA was developed and when most operators adopted it is indicative that the impact of PRA on safety and operations was not obvious. Although the primary purpose of PRA is to quantify risks at nuclear plants, we found that the process of identifying and quantifying these risks led to other benefits beyond obtaining these estimates. Even though there could be limitations to the precision of estimating risks (Rae et al., 2014), the process of measuring them has its own benefits. Moreover, we find that PRA adoption is associated with faster rates of reducing the number of significant events compared to periods before adoption.

We find evidence that the benefits of PRA extend beyond providing a numerical value of risk. Several case studies describe how operators at different plants use PRA to identify opportunities for improvements. These studies show how PRA can improve resource allocation for maintenance, broaden the operator’s knowledge of in-service repairs, and facilitate safety and operational standards (Vo & Edwards, 1994). Despite these case studies, the enthusiasm for PRA remains mixed. Our results suggest that PRA can have a large effect for poor-performing plants. PRA also provides an opportunity for operators to review industry-wide experience and learn from it.

The impact of PRA is larger at 27% (Model (6) in Table III) for operators that conducted PRA

before it was required by the NRC (proactive adopters) compared to the average decrease of 8% (Model (3) in Table III) for those that did it for the first time in response to Generic Letter 88-20 (reactive adopters). There could be several mechanisms for this observation, but we do not have enough data to disentangle which one fits best. We describe potential mechanisms next.

There are several reasons why proactive adopters experience larger benefits than reactive adopters. One possible reason is that it may take several years before most of the benefits of PRA can be clearly measured. Another possible reason is that managers who adopt early tend to do so for efficiency gains and late adopters follow to conform to external pressures (Vo, Harris, & Gore, 1997). This difference in motivation to adopt may lead to differences in outcomes. Early adopters are likely to be more deliberate, motivated in implementing new management practices to improve current ones compared to late adopters (Gray, Anand, & Roth, 2015). Managers that adopt later likely do so in response to external pressures; therefore, their benefits may be limited (Gray et al., 2015). A third possible explanation is that some of the knowledge (or “lessons learned”) from doing PRA may have already spilled over to reactive adopters before they conducted their own risk assessments. Operators and engineers can disseminate their knowledge at conferences or write about it in technical papers. For example, the idea that shorter maintenance windows can help reduce risk is knowledge that is easily transferable

to managers that have not yet conducted PRA. Although we cannot identify which mechanism best describes our observations, we do know that PRA still had a significant impact for both samples. Those that adopted PRA (reactive adopters) in response to regulation also experienced safety improvements. This is strong evidence that PRA is indeed effective even if operators simply do it in order to fulfill the request of the NRC to perform individual plant examinations.

PRA may have played a role in reducing recurring events. Conducting PRA provided operators an opportunity to think of cost-effective ways to manage possible failures either through better maintenance schedules or training. We find that the average rate of recurring events decreased from 4.3% in periods before submitting PRA to an average of 1.9% after. Although recurring events seemed low to begin with, this decrease is statistically significant ($p < 0.01$).

Our estimates suggest that conducting PRA is cost effective. A report by the U.S. General Accounting Office estimates that the cost of conducting an individual plant PRA is between \$460,000 and \$1.4 million (adjusted for inflation to reflect its value in 2018). These values may not be representative of the cost of PRA today, but they show the cost effectiveness of PRA in its early development in the industry. The reduction in the frequency of safety-related disruptions is associated with a \$13.1 million increase in annual revenue from avoided lost production. To estimate this, we merged data¹³ on monthly capacity factor, the ratio of actual electricity produced divided by the maximum possible for that month, with our data set. We then estimated the association of the number of safety-related events on the capacity factor. We found that one event is associated with a 7.4% decrease in the capacity factor for that month (see the Appendix for the regression results). This means that a 0.30 reduction in the number of safety-related events is an additional 16.4 hours (2.2% increase) of monthly electricity production. This increase is equivalent to 16,400 MWh per month for a 1,000 MW plant. At an average electricity price of \$80 per MWh, nuclear power operators can gain up to \$13.1 million per year from avoided safety-related disruptions. Overall, the direct and indirect benefits of PRA appear to easily outweigh the cost to implement them.

¹³This data set is from Davis and Wolfram (2012a).

6.2. Limitations

There are limitations to our approach. The results may not necessarily be generalizable to other types of risk assessments or in other industries. There are different types of risk assessments, but other risk assessment tools share many common features with PRA. For example, the need to collect data and identify points of vulnerability is a common feature in many risk assessment tools. There are other industries, such as the National Aeronautics and Space Administration, that use PRA. The effect of PRA will undoubtedly vary depending on the industry, but given that nuclear power is heavily regulated, the impact of PRA in other industries with less regulatory oversight could be even higher.

If our spillover theory for proactive and reactive adopters is true, then that means that the actual benefits of risk assessments are likely to be greater than what we estimate because we only estimate the same-plant benefits, not the cross-sectional benefits.

It may be the case that the magnitude of our estimates will vary depending on the industry, but given that the nuclear power is heavily regulated, our results on the impact of PRA could be underestimated compared to its impact in other industries with less regulatory oversight.

6.3. Future Work

We decided to focus on the one-time submission of PRA, but the application of PRA has evolved over time and across plants with varying levels of how integrated it is for each plant. Some plants, such as the South Texas Nuclear Generating Station, have so-called living PRA, where they continuously perform quantitative risk assessments. The discussion on the value of risk assessments and how we should conduct them remains active (Pasman, Rogers, & Mannan, 2017). We have demonstrated how the impact of PRA can be measured at the industry level with field data to supplement analytical models. Future work can explore other benefits of living PRA that may emerge beyond providing a numerical value for risk.

ACKNOWLEDGMENTS

We want to thank Dr. Warner North and the two anonymous reviewers for taking the time to review the article. The two anonymous reviewers were exceptionally knowledgeable and constructive. We want to thank Adam Fremeth, David Lochbaum,

Paulina Jaramillo, and Michael Buchdahl Roth for sharing their data on penalties, refueling schedules, and production cost. We thank Bob Budnitz for connecting us with various experts and industry members in the U.S. nuclear industry. We thank all those who shared their time and experience in the nuclear industry through interviews and e-mail exchanges: Valerie Barnes, Diane Screnci, Vicki Bier, Kenneth Koves, George Apostolakis, Regis Matzie, Andrew Kadak, Carl “Rick” Grantom, Donnie Harrison, Sheehan Neil, Tunc Aldemir, and Curtis Smith.

APPENDIX

We provide results on the association of safety-related disruptions on capacity factors.

We measure the impact of safety-related disruptions on capacity factor, the total electricity produced divided by the maximum electricity that can be generated within the same period.¹⁴ We merge monthly capacity factor data from Davis and Wolfram (2012b) with our data set. Table A1 shows the results of fixed effects models where the dependent variable is the monthly capacity factor and the key independent variable is the number of safety-related events for that month. Model (1) shows that one event is associated with a 0.074 ($p < 0.01$) decrease in capacity factor for that month. This association is robust to the inclusion of time trends in Model (2) and when weighted by plant capacity in Model (3).

Table A1. What Is the Association Between the Number of Monthly Events and Capacity Factors?

	Dependent Variable: Monthly Capacity Factor		
	(1)	(2)	(3)
Number of safety-related events	-0.074*** (0.002)	-0.072*** (0.002)	-0.071*** (0.002)
Reactor fixed effects	Yes	Yes	Yes
Time trend	-	Yes	Yes
Weighted by plant capacity	-	-	Yes
Observations	16,066	16,066	16,066
R ²	0.207	0.208	0.208

Note: Cluster-robust standard errors in parentheses.
* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

¹⁴A capacity factor of 100% means that the power plant produced the maximum electricity it can generate over that time period.

Table A2. The Association of Different Types of Events on Capacity Factors

	Dependent Variable: Monthly Capacity Factor		
	(1)	(2)	(3)
System actuation	-0.080*** (0.003)	-0.074*** (0.003)	-0.071*** (0.003)
Technical specification	-0.053*** (0.003)	-0.052*** (0.003)	-0.052*** (0.003)
Degradation	-0.094*** (0.006)	-0.099*** (0.006)	-0.102*** (0.006)
Inoperable	-0.055*** (0.014)	-0.051*** (0.014)	-0.050*** (0.014)
Prevent safety equipment	-0.033*** (0.007)	-0.031*** (0.007)	-0.027*** (0.007)
Reactor fixed effects	Yes	Yes	Yes
Time trend	-	Yes	Yes
Weighted by plant capacity	-	-	Yes
Observations	16,066	16,066	16,066
R ²	0.195	0.197	0.198

* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

Next, we examine whether the type of event may impact capacity factors differently. Table A2 shows the results where the number of safety-related events is broken down into five different types. These events are the same ones we discuss in the main article, and their description is the same here. We find that events associated with the degradation of equipment have the highest impact on capacity factor at -0.09 ($p < 0.01$), followed by system actuation at -0.08 . These two types of events are often associated with a shutdown of the plant and thus have the highest impact on production. We find that events associated with a technical specification or inoperable equipment have a -0.05 and -0.06 impact on capacity factor for that month ($p < 0.01$). Events that may have prevented a safety equipment from performing has the lowest impact on capacity factor at -0.03 , but is still significant at the 0.01 level. Although the impact may vary by type, the results show that safety-related events are negatively associated with electricity production.

REFERENCES

Apostolakis, G. E. (2004). How useful is quantitative risk assessment? *Risk Analysis*, 24(3), 515–520.

- Azizi, M. M. (2014). PRA application to offshore drilling critical systems. *Probabilistic Safety Assessment and Management*, 12, 1–13.
- Bier, V. M. (1999). Challenges to the acceptance of probabilistic risk analysis. *Risk Analysis*, 19(4), 703–710.
- Carroll, J., Sterman, J., & Marcus, A. (1998). Playing the maintenance game: How mental models drive organizational decisions. In R. N. Stern & J. J. Halpern (Eds.), *Debating rationality: Nonrational elements of organizational decision making* (pp. 99–121). Ithaca, NY: Cornell University ILR Press.
- Cooke, R. M., Ross, H. L., & Stern, A. (2011). *Precursor analysis for offshore oil and gas drilling: From prescriptive to risk-informed regulation*. Technical report, Resources for the Future.
- Corbett, C. J., Montes-Sancho, M. J., & Kirsch, D. A. (2005). The financial impact of ISO 9000 certification in the United States: An empirical analysis. *Management Science*, 51(7), 1046–1059.
- Daling, P. M., Marler, J. E., Vo, T. V., Phan, H. K., & Friley, J. R. (1995). Assessment of costs and benefits associated with resolution of generic safety issue 143—Availability of heating, ventilation, and air conditioning and chilled water systems. *Nuclear Technology*, 109(3), 429–436.
- Davis, L. W., & Wolfram, C. (2012a). Deregulation, consolidation, and efficiency: Evidence from US nuclear power. *American Economic Journal: Applied Economics*, 4(4), 194–225.
- Davis, L. W., & Wolfram, C. (2012b). Deregulation, consolidation, and efficiency: Evidence from US nuclear power: Dataset. *American Economic Journal: Applied Economics*, 4(4), 194–225.
- Dubord, R. M., Golay, M. W., & Rasmussen, N. C. (1996). A probabilistic risk assessment-related methodology to support performance-based regulation within the nuclear power industry. *Nuclear Technology*, 114(2), 169–178.
- Fremeth, A. R., & Holburn, G. L. (2012). Information asymmetries and regulatory decision costs: An analysis of US electric utility rate changes 1980–2000. *Journal of Law, Economics, and Organization*, 28(1), 127–162.
- Garrick, B. J. (1989). Lessons learned from 21 nuclear plant probabilistic risk assessments. *Nuclear Technology*, 84(3), 319–330.
- Goble, R., & Bier, V. M. (2013). Risk assessment can be a game-changing information technology: But too often it isn't. *Risk Analysis*, 33(11), 1942–1951.
- Goerlandt, F., Khakzad, N., & Reniers, G. (2016). Validity and validation of safety-related quantitative risk analysis: A review. *Safety Science*, 84, 67–77.
- Gray, J. V., Anand, G., & Roth, A. V. (2015). The influence of ISO 9000 certification on process compliance. *Production and Operations Management*, 24(3), 369–382.
- Greene, W. H. (2012). *Econometric analysis* (7th ed.). Boston, MA: Pearson.
- Hendricks, K. B., & Singhal, V. R. (1997). Does implementing an effective TQM program actually improve operating performance? Empirical evidence from firms that have won quality awards. *Management Science*, 43(9), 1258–1274.
- Keller, W., & Modarres, M. (2005). A historical overview of probabilistic risk assessment development and its use in the nuclear power industry: A tribute to the late Professor Norman Carl Rasmussen. *Reliability Engineering & System Safety*, 89(3), 271–285.
- Lochbaum, D. A. (2000). *Nuclear plant risk studies: Failing the grade*. Cambridge, MA: Union of Concerned Scientists.
- MacAvoy, P. W., & Rosenthal, J. W. (2005). *Corporate profit and nuclear safety: Strategy at northeast utilities in the 1990s*. Princeton, NJ: Princeton University Press.
- National Commission on the BP Deepwater Horizon Oil Spill (2011). *Deep water: The Gulf oil disaster and the future of offshore drilling*. Washington, DC: Perseus Distribution Digital.
- NRC Report (2018). *Background and staff guidance on performance-based regulation*. Retrieved from <https://www.nrc.gov/about-nrc/regulatory/risk-informed/concept/performance.html>.
- NRC Report (1988). *Individual plant examinations for severe accident vulnerabilities* (Generic Letter 88-20). Washington, DC: US NRC.
- NRC Report (1984). *Licensee event report system, description of system and guidelines for reporting* (Technical Report NUREG-1022). Washington, DC: US NRC.
- NRC Report (1984). *Probabilistic risk assessments (PRA): Status report guidance for regulatory application* (Technical Report NUREG-1050). Washington, DC: US NRC.
- NRC Report (1991). *Severe accident risks: An assessment for five US nuclear power plants* (Technical Report NUREG-1150). Washington, DC: US NRC.
- NRC Report (1997). *Individual plant examination program: Perspectives on reactor safety and plant performance*. (Technical Report NUREG-1560). Washington, DC: US NRC.
- NRC Report (1983). *A guide to the performance of probabilistic risk assessments for nuclear power plants* (Technical Report NUREG/CR-2300). Washington, DC: US NRC.
- NRC Report (1999). *Rates of initiating events at us nuclear power plants* (Technical Report NUREG/CR-5750). Washington, DC: US NRC.
- NRC Report (2003). *US nuclear power plant operating cost and experience summaries* (Technical Report NUREG/CR-6577). Washington, DC: US NRC.
- NRC Report (2003). *Handbook of parameter estimation for probabilistic risk assessment* (Technical Report NUREG/CR-6823, ACC: MOL. 20060126.0121) Washington, DC: US NRC.
- NRC Report (1975). *Reactor safety study, an assessment of accidents risks in US commercial nuclear power plants* (Technical Report WASH-1400). Washington, DC: US NRC.
- Pasman, H. J., Rogers, W. J., & Mannan, M. S. (2017). Risk assessment: What is it worth? Shall we just do away with it, or can it do a better job? *Safety Science*, 99, 140–155.
- Paté-Cornell, E. (1995). Managing fire risk onboard offshore platforms: Lessons from Piper Alpha and probabilistic assessment of risk reduction measures. *Fire Technology*, 31(2), 99–119.
- Paté-Cornell, M. E. (1990). Organizational aspects of engineering system safety: The case of offshore platforms. *Science*, 250(4985), 1210–1217.
- Paté-Cornell, M. E., & Bea, R. G. (1992). Management errors and system reliability: A probabilistic approach and application to offshore platforms. *Risk Analysis*, 12(1), 1–18.
- Pence, J., Abolhelm, M., Mohaghegh, Z., Reihani, S., Ertem, M., & Kee, E. (2018). Methodology to evaluate the monetary benefit of probabilistic risk assessment by modeling the net value of risk-informed applications at nuclear power plants. *Reliability Engineering & System Safety*, 175, 171–182.
- Rae, A., Alexander, R., & McDermid, J. (2014). Fixing the cracks in the crystal ball: A maturity model for quantitative risk assessment. *Reliability Engineering & System Safety*, 125, 67–81.
- Rees, J. V. (1996). *Hostages of each other: The transformation of nuclear safety since Three Mile Island*. Chicago, IL: University of Chicago Press.
- Taylor, J. B., & Wolak, F. A. (2012). *A comparison of government regulation of risk in the financial services and nuclear power industries* (Discussion Papers 11-018). Stanford: Stanford Institute for Economic Policy Research.
- U.S. General Accounting Office (1985). *Probabilistic risk assessment: An emerging aid to nuclear power plant safety regulation*.
- Vo, T. V., & Edwards, D. R. (1994). Development of in-service inspection priorities for pressurized water reactor high-pressure injection system components. *Nuclear Technology*, 106(1), 110–124.
- Vo, T. V., Harris, M. S., & Gore, B. F. (1989). Probabilistic risk assessment based inspection guidance for Arkansas Nuclear One Unit 1. *Nuclear Technology*, 84(1), 14–22.

- Vo, T. V., Simonen, F. A., Doctor, S. R., Smith, B. W., & Gore, B. F. (1993). Development of in-service inspection plans for nuclear components at the Surry unit 1 nuclear power station. *Nuclear Technology*, *102*(3), 403–415.
- Westphal, J. D., Gulati, R., & Shortell, S. M. (1997). Customization or conformity? An institutional and network perspective on the content and consequences of TQM adoption. *Administrative Science Quarterly*, *42* 366–394.
- Wheatley, S., Sovacool, B., & Sornette, D. (2017). Of disasters and dragon kings: A statistical analysis of nuclear power incidents and accidents. *Risk Analysis*, *37*(1), 99–115.
- Worledge, D. H., & Wall, I. B. (1989). Overview of the Electric Power Research Institute research program on common-cause failures. *Nuclear Technology*, *84*(3), 256–259.