SCIENCE & GLOBAL SECURITY
2016, VOL. 24, NO. 1, 37–62
http://dx.doi.org/10.1080/08929882.2016.1127039

Routledge
Taylor & Francis Group

Estimating the Frequency of Nuclear Accidents

Suvrat Raju

International Centre for Theoretical Sciences, Tata Institute of Fundamental Research, Indian Institute of Science Campus, Bengaluru, India

ABSTRACT
Bayesian methods are used to compare the predictions of probabilistic risk assessment—the theoretical tool used by the nuclear industry to predict the frequency of nuclear accidents—with empirical data. The existing record of accidents with some simplifying assumptions regarding their probability distribution is sufficient to rule out the validity of the industry's analyses at a very high confidence level. This conclusion is shown to be robust against any reasonable assumed variation of safety standards over time, and across regions. The debate on nuclear liability indicates that the industry has independently arrived at this conclusion. Paying special attention to the case of India, the article shows that the existing operating experience provides insufficient data to make any reliable claims about the safety of future reactors. Finally, policy implications of the article findings are briefly discussed.

ARTICLE HISTORY
Received 15 May 2015
Accepted 24 July 2015

Introduction

This article was motivated by the recent public debates on the safety of nuclear reactors in India. Although this is an old international debate, it has occupied a rather prominent public position in India since the passage of the Indo–U.S. nuclear deal. What is significant is that the debate has not been confined to the technical community, but has seen the active participation of peoples' movements and civil society groups. This work attempts to parse some of the claims made by the Indian nuclear establishment and international nuclear vendors in this debate. Of course, while the Indian discussion provided motivation and also specialized some of our findings to India, the central conclusion about the conflict between the results of probabilistic risk assessment and empirical data is valid more broadly.

This article reviews and analyzes the "probabilistic risk assessment" (PRA) framework that is used by the nuclear industry to calculate the expected frequency of nuclear accidents. This framework often produces expected frequency that is extremely low.

CONTACT Suvrat Raju suvrat.raju@gmail.com TIFR Centre Building, Indian Institute of Science Campus, Bengaluru, 560012, India.
Color versions of one or more of the figures in the article can be found online at www.tandfonline.com/gsgs.

Although it is likely that the industry and policy-makers are privately aware of the unreliability of these figures, they are nevertheless used quite commonly in policy debates. Thus, for example, the previous chairperson of the Indian Atomic Energy Commission declared that the chance of a nuclear accident in India was "1-in-infinity."[1]

Similar statements can be found in the scholarly literature. For example, in a review of the 1000 MW VVER reactors that were installed in the South Indian city of Kudankulam, officials from the Nuclear Power Corporation of India (NPCIL) pointed out that the Core-Damage Frequency (CDF) of these reactors—the rate at which the reactor is expected to suffer accidents that damage its core—was just 10^{-7} per reactor-year (ry).[2]

The other two multinational companies that are in line to construct reactors in India—the French company, Areva, whose European Pressurised Reactors (EPRs) have been selected for Jaitapur (Maharashtra) and the the American company Westinghouse, whose AP1000 reactors have been selected for Mithi Virdi (Gujarat)—have also put forward similar figures. For example, Areva claims that the CDF of the EPR is $7.08 \times 10^{-7}(\mathrm{ry})^{-1}$.[3] Westinghouse has estimated the CDF of the AP1000 to be $5.09 \times 10^{-7}(\mathrm{ry})^{-1}$.[4]

Prima facie, it is prudent to treat such low numbers skeptically, especially when they are provided with a precision of two decimal places. A nuclear reactor is a very complex system, and our understanding of its dynamics, particularly when they are coupled to an uncertain external environment, is not advanced enough to permit such accurate estimates.

Nevertheless, these numbers are taken seriously by regulators, both in India and elsewhere. In the United States, the Nuclear Regulatory Commission has set a goal for PRA estimates of both the CDF and the large-release frequency (LRF). The latter involves accidents where the containment fails in addition to core damage. These are expected to be below $10^{-4}(\mathrm{ry})^{-1}$ and $10^{-6}(\mathrm{ry})^{-1}$ respectively.[5,6]) Similar quantitative criteria have been adopted in other countries.[7] In India, licensing guidelines suggest that new plants should have a PRA-estimated CDF smaller than $10^{-5}(\mathrm{ry})^{-1}$ and a LRF lower than $10^{-6}(\mathrm{ry})^{-1}$.[8]

The methodology used to carry out PRAs in the nuclear industry is theoretically suspect. However, the purpose of this article is to point out that, even setting aside theoretical considerations, these extraordinarily low bounds can be ruled out by using the existing empirical record of nuclear accidents.

As stressed in this article, there have already been eight core-damage accidents in a little more than fifteen thousand reactor-years of experience. Hence, the observed frequency of accidents is significantly higher than that predicted by the industry's PRAs. Moreover, and this is the crucial conclusion of this article: had the true frequency of accidents been as low as the manufacturers claim, it is exceedingly unlikely that so many accidents could have occurred. Conversely, the historical record on accidents implies that, even under rather conservative assumptions, it is possible to conclude with a very high degree of confidence that the results of the industry's PRAs are incorrect.

For example, the article shows that in a simplified model, even with favorable assumptions for the industry, the hypothesis that the frequency of core-damage accidents is smaller than or equal to $10^{-7}(\text{ry})^{-1}$ can be ruled out with a confidence of $1.0 - 4.0 \times 10^{-24}$. (Said differently, within this model, it can be concluded that the core-damage frequency of nuclear reactors is higher than $10^{-7}(\text{ry})^{-1}$ with a confidence of 99.9999999999999999999996%!) The precise value of 4.0×10^{-24} is specific to this model but the broader and robust conclusion is that the accident frequencies suggested by PRA can be ruled out, even with the limited empirical data on accidents, at an extremely high confidence level. Using a sensitivity analysis, it is shown that even by allowing exponential increases in safety standards over time, or significant variations across regions, it is virtually impossible to reconcile the empirical data with the PRA-frequencies.

Finally, the article turns briefly to the specific case of India and makes the elementary statistical point that the existing experience of Indian reactor operation is insufficient to derive any strong conclusions about expected accident frequencies in the future. This result is then discussed from a policy standpoint, and placed in perspective with the nuclear industry position in the debate on the liability of nuclear accidents.

Summary and methodology

The central objective of this article is to use the empirical data to estimate the probability for the hypothesis that the true frequency of accidents is indeed as low as that claimed by PRA. This probability is denoted by $\mathcal{C}(\lambda_{\text{pra}}, n_{\text{obs}})$ where λ_{pra} is a frequency predicted by PRA and n_{obs} is the observed number of accidents of a certain type. The central result is that

$$\mathcal{C}(\lambda_{\text{pra}}, n_{\text{obs}}) = \epsilon \ll 1. \tag{1}$$

The precise value of ϵ depends on many factors: the precise PRA being considered, which leads to variations in λ_{PRA}, the way in which accidents are counted which leads to variations in n_{obs} and, of course, the assumptions that go into modeling nuclear accidents and the variation of safety-standards over time and across regions.

But the point emphasized in this article is the following: ϵ is extraordinarily small, and this fact is robust against virtually any variation of the considerations above. The calculations presented in this article are all calculations of ϵ under different assumptions and they all have the feature that ϵ is negligible.

It is important to note that this article does not intend to compute the true frequency of accidents using empirical data. Some efforts have been made in this direction,[9,10] but the basic difficulty is that the empirical data is too limited to allow for a reliable computation of this frequency. The importance of focusing on ϵ is that these intermediate uncertainties do not affect the robust nature of equation (1).

Second, note that to compute ϵ, Bayesian analysis must be used. While a purely frequentist approach might demonstrate that the empirically observed frequency of accidents does not agree with the frequency predicted by PRA, this does not, by itself, tell the confidence level with which the correctness of PRA can be ruled out. In fact the question of estimating the confidence-level cannot be posed within the frequentist approach at all. This confidence-level is given by $1 - \epsilon$, and requires Bayesian methods.

A brief review of probabilistic risk assessment

The systematic use of PRA in the nuclear industry started with the Rasmussen report of 1975.[11] This report was commissioned by the U.S. Nuclear Regulatory Commission (NRC) and met with criticism soon after its publication. In 1977, the U.S. NRC commissioned a review of the Rasmussen report through the Lewis Committee, and following this critical review-report, released a statement in 1978 stating that "the Commission does not regard as reliable the Reactor Safety Study's numerical estimates of the overall risk of reactor accidents."[12]

In spite of these early objections, over the past few decades, the framework of probabilistic risk assessment has become influential in propagating the use of quantitative probabilistic techniques for questions of nuclear safety. Moreover, as described above, the use of these techniques is not confined to qualitative safety analyses aided by quantitative probabilistic techniques, but rather the actual numerical results produced by PRA are used by regulators, and in policy debates.

The basic idea used by such studies is simple to describe: one enumerates the possible fault trees that could lead to an accident. For each individual component in the reactor, one can estimate a frequency of failure. For a serious accident, some combination of these components has to fail simultaneously. The industry advertises a philosophy of "defense in depth," which reduces the overall possibility of an accident by building redundancy into a system. The low numbers above result from the fact that several systems have to fail simultaneously before the core is damaged.

For example, Keller and Modarres report that in 1966,[13] in one of the earliest such assessments, the General Electric company "showed" that its reactors "had a one-in-a-million chance per year for a catastrophic failure because each of the three major subsystems would only fail once-in-one-hundred per years" *(sic)*. The Rasmussen report, which followed a decade later, attempted to refine and formalize this methodology.

The theoretical problem with such estimates is obvious. Consider the Fukushima nuclear complex, which had 13 backup diesel generators.[14] Assigning a probability of 10^{-1} for the failure of each generator per year and assuming that they are independent would lead to the naive conclusion that the probability that 12 generators would fail together in any given year is about $13 \times 10^{-12} \times 0.9 \approx 10^{-11}$. However, the tsunami did precisely this by disabling all but one of the generators at once. The point is that once the obvious fault trees have been eliminated and corrected, the dominant contributions to accident-probabilities come from unlikely sequences of events that conspire to cause failure.

Table 1. Predictions of Probabilistic Risk Assessment.

Reactor (Manufacturer)	Core-Damage Frequency	Large-Release Frequency
EPR (Areva)	$7.08 \times 10^{-7}(\text{ry})^{-1}$	$7.69 \times 10^{-8}(\text{ry})^{-1}$
AP1000 (Westinghouse)	$5.09 \times 10^{-7}(\text{ry})^{-1}$	$5.94 \times 10^{-8}(\text{ry})^{-1}$
VVER (Rosatom)	$10^{-7}(\text{ry})^{-1}$	$10^{-8}(\text{ry})^{-1}$

This issue occurs in any suitably complex system. However, a nuclear reactor is also an open system that is coupled to an external environment. This makes it virtually impossible to foresee all sorts of low probability pathways to failure. Furthermore, the understanding of the frequency of extreme initiating events, such as tsunamis and earthquakes, is itself rather crude. These error bars overwhelm the seemingly accurate predictions that result from elaborate simulations of the reactor.

However, instead of venturing deeper into these theoretical arguments—which are considered in greater detail elsewhere in the literature,[15]—the goal is to examine how the predictions of PRA stand up against the extant empirical data.

To facilitate this comparison, the various kinds of claims that have been made by the three multinational companies relevant to India are summarized. As mentioned, Areva has estimated that the CDF of an EPR accounting for both internal and external hazards is $7.08 \times 10^{-7}(\text{ry})^{-1}$. The LRF of the EPR is estimated to be about 11% of its CDF: $7.69 \times 10^{-8}(\text{ry})^{-1}$.[3,16]

Similarly, Westinghouse has estimated that the CDF of the AP1000 is $5.09 \times 10^{-7}(\text{ry})^{-1}$. The LRF of this reactor is estimated to be $5.94 \times 10^{-8}(\text{ry})^{-1}$.[4] This includes accidents due to "external hazards" including "external flooding, extreme winds, seismic, and transportation accidents." In fact Westinghouse concluded that "conservative bounding assessments show that core damage risk from events listed above is small compared to the core damage risk from at-power and shutdown events."[4]

Although the NPCIL has stated that the CDF of the Kudankulam VVER reactors is $10^{-7}(\text{ry})^{-1}$, the details of the PRA that led to this conclusion could not be located. So, this figure will be assumed as is, giving a LRF of $10^{-8}(\text{ry})^{-1}$ using the common estimate that the LRF "is generally about ten times less than CDF."[17] All these claims are summarized in Table 1.

Review of the empirical experience

In this section, the historical record on nuclear accidents is reviewed. The industry, as a whole, had gathered about $T_{\text{obs}} = 15,247$ reactor-years of operating experience by the end of 2012 according to the latest data put out by the International Atomic Energy Agency (IAEA).[18] In this time, there have been several core-damage accidents.

Surprisingly, the IAEA does not maintain a comprehensive historical record of core-damage accidents. Cochran and McKinzie have compiled a very useful list of 25 such instances from various sources.[19] This analysis considers only accidents that occurred at commercial reactors, thus excluding accidents at experimental facilities

Table 2. List of core-damage accidents in commercial reactors.

Reactor	Country	Year	Note
Saint-Laurent A-1	France	1969	Meltdown of 50 kg of fuel.[20]
Three Mile Island Unit 2	USA	1977	Severe accident; radiation release.[21]
Saint-Laurent A-2	France	1980	Meltdown of one channel of fuel.[19,22]
Chernobyl Unit 4	Ukraine	1986	Severe accident; large radiation release.[23]
Greifswald Unit 5	Germany (GDR)	1989	Partial core meltdown soon after commissioning.[19, 24-26]
Fukushima Daiichi Units 1,2,3	Japan	2011	Severe accident; large radiation release.[27]

like Enrico Fermi Unit-1 (1966) or Lucens (1969). Even this enumeration is somewhat subjective, but a conservative approach, keeping only accidents that involved a significant meltdown of fuel leads to the list in Table 2. In each case, references to more detailed descriptions of these accidents are provided.

Table 2 enumerates accidents at $n^{\rm cd}_{\rm obs} = 8$ reactors. Of these 8 accidents, $n^{\rm lr}_{\rm obs} = 5$ accidents led to the release of large amounts of radioactive substances into the environment. This list comprises the accidents at Three Mile Island, Chernobyl, and the three at Fukushima. With the available PRA results for both the CDF and the LRF, it is possible to compare these separately to the historical record.

The reason for counting the three accidents at Fukushima separately in the list above is to calculate the rate of accidents per reactor-year of operation, since all these three reactors contribute separately to the "total operating experience" that appears in the denominator.

However, to demonstrate the robust nature of the conclusions of this article, a parallel analysis where the accidents at Fukushima are counted together is also presented. By this *incorrect* counting, there have been $n^{\rm cd}_{\rm low} = 6$ core damage accidents and $n^{\rm lr}_{\rm low} = 3$ accidents with a large release of radioactivity.

Simplified Bayesian analysis of empirical frequencies and PRA

Table 2 leads to the following observed frequency of core-damage and large-release accidents:

$$\nu^{\rm cd}_{\rm obs} = \frac{n^{\rm cd}_{\rm obs}}{T_{\rm obs}} \approx \frac{1}{1906}({\rm ry})^{-1} \approx 5.2 \times 10^{-4}({\rm ry})^{-1},$$
$$\nu^{\rm lr}_{\rm obs} = \frac{n^{\rm lr}_{\rm obs}}{T_{\rm obs}} \approx \frac{1}{3049}({\rm ry})^{-1} \approx 3.3 \times 10^{-4}({\rm ry})^{-1}. \quad (2)$$

Similarly, counting the Fukushima accidents as a single accident,

$$\nu^{\rm cd}_{\rm low} = \frac{n^{\rm cd}_{\rm low}}{T_{\rm obs}} \approx \frac{1}{2541}({\rm ry})^{-1} \approx 3.9 \times 10^{-4}({\rm ry})^{-1},$$
$$\nu^{\rm lr}_{\rm low} = \frac{n^{\rm lr}_{\rm low}}{T_{\rm obs}} \approx \frac{1}{5082}({\rm ry})^{-1} \approx 2.0 \times 10^{-4}({\rm ry})^{-1}.$$

It is clear that these empirically observed frequencies are far higher than the predictions of the manufacturers' PRAs. However, a more detailed question can be asked: given the observed rate of accidents, what is the probability that the results of PRA are close to, or smaller than, the "true frequency" of accidents. Bayesian techniques are ideally suited to answer this question.

To answer the question above, a few simplifying assumptions about the probability distribution of nuclear accidents are needed. As a first approximation, it can be assumed that nuclear accidents are independent events, and that in every small time interval dt, each reactor has a small and constant probability

$$dp = \lambda\, dt, \tag{3}$$

of suffering an accident. The use of this approximation is *not meant* to suggest that it is truly the case that accident frequencies have not changed over time. Rather, as it is shown later, the results are very *robust* against almost any reasonable assumed changes in safety standards over time, and across regions.

Therefore, given this robustness of the central results, a simplified model is presented first since it is amenable to simple analysis and already captures the central points of this article. A more sophisticated analysis is presented in the next section for completeness.

Assuming that m reactors are functioning simultaneously and are observed for a time period $T_1 = Ndt$. The probability for n of these to have undergone accidents is

$$\begin{aligned} p_\lambda(n) &= \binom{m}{n}(1-\lambda dt)^{N(m-n)}\left(1-\left(1-\lambda dt\right)^N\right)^n \\ &= \binom{m}{n}\left(1-\frac{\lambda T_1}{N}\right)^{N(m-n)}\left(1-\left(1-\frac{\lambda T_1}{N}\right)^N\right)^n \\ &\underset{N\to\infty}{\longrightarrow} \frac{\Gamma(m+1)}{\Gamma(m-n+1)\Gamma(n+1)} e^{-\lambda T_1 m(1-\frac{n}{m})}\left(1-e^{-\lambda T_1}\right)^n, \end{aligned}$$

where, the last line is the continuous limit for $N \to \infty$, with T_1 finite and $\Gamma(n+1) = n!$ is the standard Gamma function. Considering the case where $\lambda T_1 \ll 1$, which implies that the chance of any individual reactor undergoing an accident is small, and $n \ll m$, which states that only a small fraction of all reactors undergo an accident, the total operating experience gathered becomes (in the continuous limit), $T = mT_1$, and the distribution simplifies to

$$p_\lambda(n) = \frac{1}{\Gamma(n+1)}(\lambda T)^n e^{-\lambda T}. \tag{4}$$

This is simply a Poisson distribution, which is commonly used to model accidents and other rare events in various scenarios.

It is now possible to solve the following Bayesian problem: start with the prior assumption that λ is uniformly distributed

$$\mathcal{P}(\lambda) = \frac{\theta(\lambda) - \theta(\lambda - \lambda^c)}{\lambda^c}, \tag{5}$$

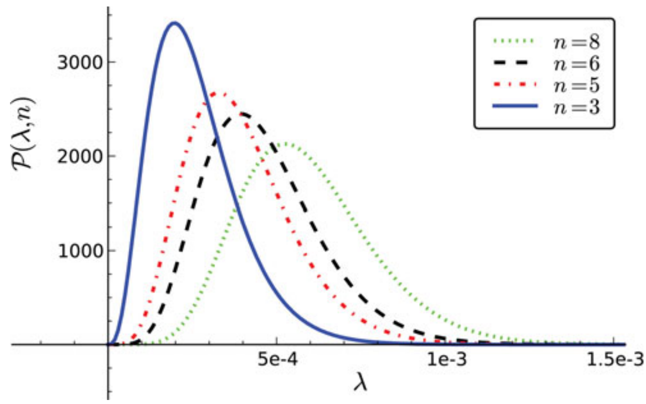

Figure 1. Posterior probability distribution for the parameter λ.

where λ^c is an irrelevant high frequency cutoff to make the distribution normalizable. Given the observed frequency of events above, what is the posterior probability distribution for λ?

Or in simpler words, starting with no prior bias for the value of λ and given that there exists some number n_{obs} of accidents, the empirical data can be used to form an estimate of the value of λ. In fact, Bayes theorem allows the actual calculation of the probability that the true frequency of accidents has a value between λ and $\lambda + d\lambda$ through the formula

$$\mathcal{P}(\lambda|n_{\text{obs}}) = \frac{\mathcal{P}(n_{\text{obs}}|\lambda)\mathcal{P}(\lambda)}{\mathcal{P}(n_{\text{obs}})}.$$

This formula makes precise the intuition that the empirical evidence is providing some information about the true value of λ. On the right side, $\mathcal{P}(n_{\text{obs}}|\lambda) = p_\lambda(n_{\text{obs}})$, which is given by the Poisson distribution Eq. 4. $\mathcal{P}(\lambda)$ is the flat prior probability distribution in Eq. 5. $\mathcal{P}(n_{\text{obs}})$, which is a λ independent constant, is fixed by $\int_0^\infty \mathcal{P}(\lambda|n_{\text{obs}})d\lambda = 1$.

After fixing this constant, the posterior probability distribution for λ is given by

$$\mathcal{P}(\lambda|n_{\text{obs}}) = \frac{1}{\Gamma(n_{\text{obs}}+1)} T_{\text{obs}}(\lambda T_{\text{obs}})^{n_{\text{obs}}} e^{-\lambda T_{\text{obs}}}, \tag{6}$$

where terms of $\mathrm{O}(e^{-\lambda_c T_{\text{obs}}})$ are neglected, assuming λ_c is large enough. This function is plotted in Figure 1 for the values of n_{obs} that are relevant to both the large-release and the core-damage frequency.

It is now important to mention a somewhat subtle point. Since the observed number of events is small, $n_{\text{obs}} \sim \mathrm{O}(1)$, the curves in Figure 1 have an appreciable width. This indicates the difficulty with using a *frequentist* approach to estimate the true frequency of accidents using empirical data. However, as it was emphasized before, if the interest is in estimating the parameter ϵ in Eq. 1 instead, then Bayesian methods yield a robust statement.

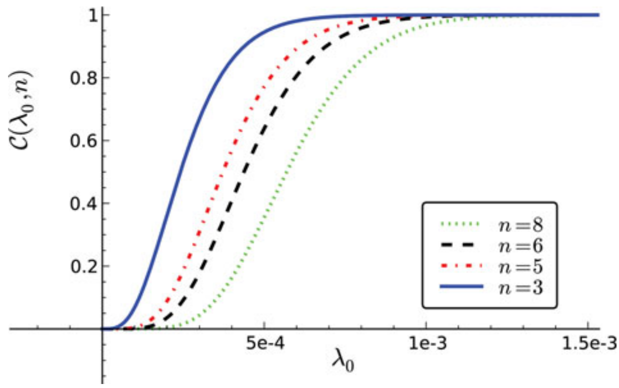

Figure 2. Probability for the hypothesis $\lambda < \lambda_0$.

To estimate ϵ, Eq. 6 is used to calculate the probability that λ is smaller than any given λ_0. This function is given by

$$\mathcal{C}(\lambda_0, n_{\text{obs}}) = \int_0^{\lambda_0} \mathcal{P}(\lambda|n_{\text{obs}})\, d\lambda = 1 - \frac{\Gamma(1 + n_{\text{obs}}, T_{\text{obs}}\lambda_0)}{\Gamma(1 + n_{\text{obs}})}. \tag{7}$$

where $\Gamma(k, z)$ is the incomplete gamma function. The function $\mathcal{C}(\lambda_0, n)$ is shown in Figure 2 for all the relevant values of n.

Since the probability that the true frequency is smaller than the various results of PRA is so close to zero, it cannot be read off the graph. However, numerical figures can be obtained using the following series expansion:

$$\Gamma(k, z) = \Gamma(k) - \frac{x^k}{k} + \frac{x^{k+1}}{k+1} + \mathrm{O}\left(x^{k+2}\right), \tag{8}$$

and are summarized in Table 3. The phrase “PRA... is right” is shorthand for the hypothesis that the true frequency of accidents is lower than or equal to the frequency predicted by PRA. Therefore, the table lists the values of the function $\mathcal{C}(\lambda_0, n_{\text{obs}})$ from Eq. 7 with λ_0 set to the PRA-predicted frequency and n_{obs} set to the observed number of accidents, which varies depending on whether the Fukushima accidents are counted together or separately. Note that the values for the probabilities in this table correspond precisely to the parameter ϵ in Eq. 1.

One observes immediately that, given the empirical data, the probability that the industry’s PRA-based conclusions are right is astronomically small. As stated in the introduction, this implies that with almost perfect certainty it can be concluded that

Table 3. Comparing PRA results with Bayesian estimates from historical observations (Simplified Model).

	Probability PRA CDF is right		Probability PRA LRF is right	
Reactor	Fukushima separate	Fukushima together	Fukushima separate	Fukushima together
Kudankulam	1×10^{-31}	4×10^{-24}	2×10^{-26}	2×10^{-17}
EPR	5×10^{-24}	3×10^{-18}	4×10^{-21}	8×10^{-14}
AP1000	3×10^{-25}	3×10^{-19}	8×10^{-22}	3×10^{-14}

the true frequency of accidents is much larger than the figures advertised by the manufacturers.

Figures in Table 3 should not be used as precise numerical bounds on the validity of the PRA results, however. The precise values of ϵ listed there suffer from several uncertainties that have already been mentioned. As shown in the next section upon consideration of a more sophisticated model, the numerical values of these probabilities change but the robust statement is that they always remain extremely small.

In words, the results of Table 3 can be stated by means of the following straightforward conclusion: the historical data on nuclear accidents provide overwhelming evidence that the methodology of probabilistic risk assessment is seriously flawed. A corollary is that the observed frequency of accidents contradicts the industry's claim that the probability of an accident is negligible.

Robustness of the simplified model

In this section, both possible improvements in safety standards over time and variations in accident frequencies across regions are modeled in the analysis. This shows that the central results of the simplified Bayesian model above are very robust against any reasonable assumed variations of this kind.

Modelling improvements in safety over time

To model possible improvements in safety over time, the assumption made in Eq. 3 are relaxed and the probability of an accident is allowed to vary with time. First, a general framework to model this possibility is discussed. Then, a concrete model is presented.

Framework for time-variations of safety

The probability that an accident occurs between time t and $t + dt$ is given by

$$dp(t) = \lambda_{\{\alpha\}}(t)dt,$$

where the subscript $\{\alpha\}$ indicates a set of parameters that control the variation of λ in time.

Now, consider a time interval of length T, which is divided into N equal parts set off by $0 < t_1 < t_2 \ldots t_{N-1} < T$ with $t_{i+1} - t_i = dt$. To specify the pattern of accidents note now that a single number is not enough, but instead requires a *function* $N(i)$, which specifies whether an accident occurred in the interval $[t_i, t_{i+1}]$. This is defined through

$$N(i) = \begin{cases} 1, \text{ if an accident occursin the interval } [t_i, t_{i+1}] \\ 0, \text{ otherwise} \end{cases}$$

Clearly, the probability for any such pattern of accidents is given by

$$p(\{N_i\}|\{\alpha\}) = \Upsilon_{\{\alpha\}} \prod_{N_i=1} \lambda_{\{\alpha\}}(t_i)dt \prod_{N_i=0} (1 - \lambda_{\{\alpha\}}(t_i)),$$

where $\Upsilon_{\{\alpha\}}$ is a normalization factor that is discussed below. Note that this probability is infinitesimal for any given pattern $N(i)$, and this is not surprising since in the continuum limit, a *path integral* over all possible patterns of accidents is needed. However, as will be shown below for the purpose of constructing the posterior probability distribution, this constant will not be important.

In the continuum limit, if n accidents are observed at times $t_{i_1} \ldots t_{i_n}$, then

$$\mathcal{P}\left(\{N_i\}|\{\alpha\}\right) \propto \left[\prod_j \lambda_{\{\alpha\}}(t_{i_j})\right] e^{-\int_0^T \lambda(t)dt}$$

Using Bayes theorem,

$$\mathcal{P}\left(\{\alpha\}|\{N_i\}\right) = \frac{\mathcal{P}\left(\{N_i\}|\{\alpha\}\right)\mathcal{P}(\{\alpha\})}{\mathcal{P}(\{N_i\})}$$

Therefore, the unknown normalization constants drop out and the *posterior probability distribution* for the parameters upon observation of a certain pattern of accidents is given by

$$\mathcal{P}\left(\{\alpha\}|\{N_i\}\right) = \mathcal{N}\left[\prod_j \lambda_{\{\alpha\}}(t_{i_j})\right] e^{-\int_0^T \lambda(t)dt}\mathcal{P}(\{\alpha\})$$

where $\mathcal{P}(\{\alpha\})$ is the *prior probability distribution* for the parameters and the normalization $\mathcal{N}$ is now a simple finite quantity that is simply fixed by

$$\int \mathcal{P}\left(\{\alpha\}|\{N_i\}\right) d\{\alpha\} = 1$$

A concrete model

This section turns to a concrete model that shows how the framework above may be utilized. First, it is assumed that λ variates with time as

$$\lambda(t) = \lambda_i e^{-\gamma t}. \tag{9}$$

Here t is the total operating experience accumulated by the industry. For $\gamma > 0$, this model proposes that the frequency of accidents starts with λ_i and decreases exponentially with time as the industry gains additional experience.

This model of an exponential increase in safety standards constitutes a very favorable assumption for the industry, but as it will be shown, is not enough to affect the central conclusions of this article. Of course, these results can easily be generalized to different variations of the frequency.

Furthermore, a flat prior distribution of the form Eq. 5 is assumed for both λ_i and γ,

$$\mathcal{P}(\lambda_i) = \frac{\theta(\lambda_i) - \theta(\lambda_i - \lambda_i^c)}{\lambda_i^c}$$
$$\mathcal{P}(\gamma) = \frac{\theta(\gamma) - \theta(\gamma - \gamma^c)}{\gamma^c}, \tag{10}$$

where λ_i^c and γ^c are cutoffs. As described above, the cutoff on λ_i is irrelevant, but in this exponential model the cutoff γ^c requires careful attention. Assuming a prior that is flat over very large ranges of γ, allows for a fat-tail in the posterior distribution that represents the scenario where the accident frequencies were very large in the past but have improved rapidly very recently. A specific numerical choice of the cutoff is therefore required.

It is important to emphasize two subtleties. First, while the cutoff, γ^c is clearly physically important and changes the numerical values of probabilities, it it not strictly required for convergence. Second, while the flat priors in Eq. 10 reflect ignorance about these parameters this necessarily involves a choice of basis. Note, for example, that assuming a flat prior for the initial frequency λ_i is different from assuming a flat prior for the current frequency $\lambda(T)$.

Now consider the case where n accidents have been observed at times $t_{i_1}, \ldots t_{i_n}$ in a total operating time T. Define $\tau = \sum_j t_{i_j}$. Then it is clear from the analysis above that the posterior probability distribution for λ_i, γ is

$$\mathcal{P}\left(\lambda_i, \gamma | \{N_i\}\right) = \mathcal{N} \lambda_i^n e^{-\gamma\tau - \frac{\lambda_i}{\gamma}\left(1 - e^{-\gamma T}\right)} \tag{11}$$

The marginal probability distributions for both λ_i and γ can be determined by integrating over the other variable. In particular, the distribution for γ can be obtained by doing the easy integral over λ_i and is given by

$$\mathcal{P}\left(\gamma | \{N_i\}\right) = \int_0^\infty d\lambda_i \mathcal{P}\left(\lambda_i, \gamma | \{N_i\}\right) = \mathcal{N} e^{-\gamma\tau} \Gamma(n+1) \left(\frac{\gamma}{1 - e^{-\gamma T}}\right)^{n+1}$$

On the other hand, it does not appear possible to write the distribution for λ_i in terms of elementary functions. However, a double-infinite series representation can be obtained as follows by expanding the exponentials.

$$\begin{aligned}
\mathcal{P}\left(\lambda_i | \{N_i\}\right) &= \int_0^{\gamma^c} d\gamma \mathcal{P}\left(\lambda_i, \gamma | \{N_i\}\right) \\
&= \mathcal{N} \int_0^{\gamma^c} d\gamma \sum_{m,q=0}^{\infty} \left(\frac{(-1)^q e^{-\frac{\lambda_i}{\alpha}} \lambda_i^{m+n} \alpha^{q-m} (mT + \tau)^q}{\Gamma(m+1)\Gamma(q+1)}\right) \\
&= \mathcal{N} \sum_{m,q=0}^{\infty} \frac{(-1)^q \lambda_i^{n+q+1} (mT + \tau)^q \Gamma\left(m - q - 1, \frac{\lambda_i}{\gamma^c}\right)}{\Gamma(m+1)\Gamma(q+1)}.
\end{aligned}$$

Table 4. Operating Experience (Reactor-Years) Accumulated by the Industry at the Time of Each Accident.

Accident	Year	Operating Experience (ry)
Saint-Laurent A-1	1969	391
Three Mile Island	1977	1406
Saint-Laurent A-2	1980	2048
Chernobyl Unit 4	1984	3150
Greifswald Unit 5	1989	5061
Fukushima Units 1,2,3	2011	14572

By integrating these distributions, the value of the normalization constant $\mathcal{N}$ is obtained through

$$1 = \int \mathcal{P}\left(\gamma|\{N_i\}\right) d\gamma = \int \mathcal{P}\left(\lambda_i|\{N_i\}\right) d\lambda_i$$
$$= \mathcal{N} \sum_{q=0}^{\infty} \frac{\Gamma(n+q+1)(qT+\tau)^{-n-2}(\Gamma(n+2) - \Gamma(n+2, \gamma^c(qT+\tau)))}{\Gamma(q+1)}$$

which yields an expression for $\mathcal{N}$ as the inverse of an infinite sum.

Another interesting quantity is the posterior probability distribution for the "current accident frequency" $\lambda_T = \lambda_i e^{-\gamma T}$. Changes can be made to the variables in Eq. 11, after including a Jacobian factor, to obtain

$$\mathcal{P}\left(\lambda_T, \gamma|\{N_i\}\right) = \mathcal{N} \lambda_T^n e^{-\gamma\tau + \frac{\lambda_T\left(1-e^{\gamma T}\right)}{\gamma} + \gamma(n+1)T}$$

As usual the probability that λ_i is smaller than a given value, λ_0, or the probability that γ is smaller than some γ_0 is given by integrating the probability distributions above.

$$\mathcal{C}_\gamma\left(\gamma_0, \{N_i\}\right) = \int_0^{\gamma_0} \mathcal{P}\left(\gamma|\{N_i\}\right) d\gamma$$
$$\mathcal{C}_i(\lambda_0, \{N_i\}) = \int_0^{\lambda_0} \mathcal{P}\left(\lambda_i|\{N_i\}\right) d\lambda_i$$
$$\mathcal{C}_T(\lambda_T^0, \{N_i\}) = \int_0^{\lambda_T^0} d\lambda_T \int_0^{\gamma^c} d\gamma \mathcal{P}\left(\lambda_T, \gamma|\{N_i\}\right)$$

While it is possible to write these expressions in terms of an infinite series of elementary functions, these forms are too complicated to be useful. It is, of course, possible to evaluate these expressions numerically at any given point as done below.

The discussion now turns to the numerical values of these expressions at the empirically relevant points. Since the expressions above depend on the specific times at which the accidents occurred, some additional information are required: the reactor-years of operating experience that the industry had accumulated at the time of the accidents listed in Table 2. Using information from the literature,[18] these figures can be estimated for the years in which the accidents occurred, by simply adding the number of operating reactors, as given in that table. This estimate is more than sufficient for our purposes and results in the Table 4.

Table 5. Comparing PRA results with Bayesian estimates from historical observations after allowing exponential increase in safety standards

	Probability PRA CDF is right		Probability PRA LRF is right	
Reactor	Fukushima separate	Fukushima together	Fukushima separate	Fukushima together
Kudankulam	7×10^{-24}	2×10^{-17}	3×10^{-22}	5×10^{-14}
EPR	3×10^{-16}	1×10^{-11}	5×10^{-17}	2×10^{-10}
AP1000	1×10^{-17}	1×10^{-12}	1×10^{-17}	6×10^{-11}

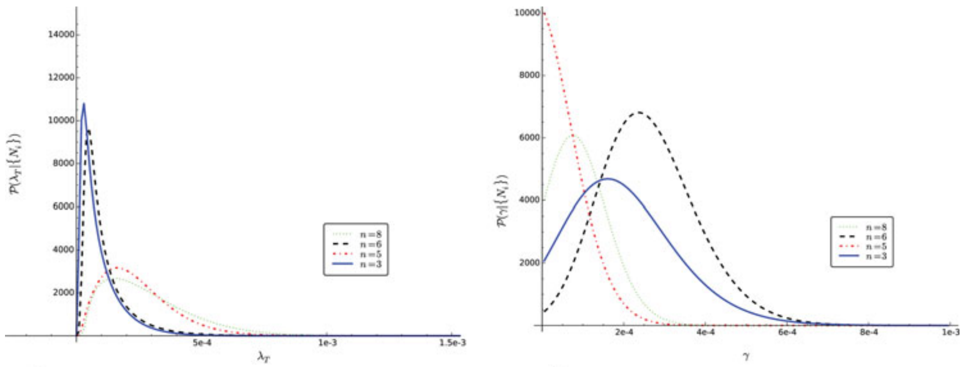

(a) Current Accident Frequency *(b) Rate of Improvement*

Figure 3. Posterior Probability Distributions for the two parameters of the model.

Finally, to obtain numerical values, the cutoff on the rate of improvement is set at

$$\gamma^c = \frac{\ln(50)}{T_{\text{obs}}}$$

which indicates the prior assumption that safety standards in the industry have improved by, at most, a factor of 50 from early commercial nuclear reactors. Of course, other values of the cutoff can be considered and, as mentioned earlier, it is even possible to remove the cutoff altogether.

With these numerical values, it is possible to plot the various probability distribution functions given above. The posterior probability distributions for λ_T and γ are plotted in Figure 3. Note that the probability distribution for $n = 5$ is peaked to the right of the distribution for $n = 6$. This is because the case with $n = 5$ represents the analysis for large-release accidents, where the Fukushima accidents are counted separately. This tends to disfavor large values of γ, since in this counting, 3 out of 5 accidents happened at late times. By thus disfavoring rapid recent improvements in safety, this particular case tends to disfavor low values of λ_T, and this is why the distribution with $n = 6$ is peaked to the right of the case with $n = 5$ even though the observed number of accidents is larger in this case.

Eventually, this analysis is interested in the probability for the hypothesis that the true frequency of accidents is less than or equal to the value predicted by PRA. Using these values, an analogue of Table 3 is obtained. Below, the relevant values of $\mathcal{C}(\lambda_T^0|\{N_i\})$ are displayed at the same points as in Table 3. This is the probability that

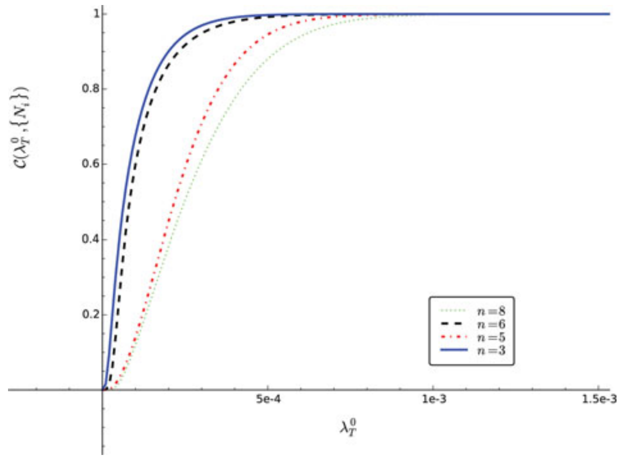

Figure 4. Probability for the hypothesis $\lambda_T < \lambda_T^0$.

the true current frequency of accidents is smaller than the values given by the PRAs considered below. This probability is also plotted as a function of λ_T^0 in Figure 4.

It is sometimes believed that since safety standards have been improving over time, the results of PRA may be valid for newer reactors even if they are inconsistent with the empirical data for older reactors. The calculations above rule out this possibility and reinforce the conclusion that the results of PRA cannot be reconciled with empirical data.

Regional variations

It is clear that the safety of nuclear reactors may vary across regions. However, just as in the analysis of possible improvements in safety over time, it can be shown that no reasonable variation in safety across regions can lend any confidence to the results of PRA-calculations.

In the case of regional variations, it would be inappropriate to proceed along the lines of Eq. 9 since there is no reason to expect that safety will vary monotonically along any spatial parameterization of nuclear reactors.

However, the central point is that for any particular country or region, the experience of nuclear accidents in the rest of the world provides a reasonable *prior* estimate of the frequency of accidents in that region. If there is reason to believe, from the record of smaller incidents or from some other information, that safety in that region is better or worse than other parts of the world, this factor can be accounted for in the prior distribution. This prior distribution can then be corrected using empirical data from the region itself to obtain a final posterior distribution for the frequency of accidents.

The discussion now describes, more precisely, how this procedure can be implemented. Consider a particular region, *R*, which may be a country or a group of countries. In this region, it is assumed that the probability of an accident in the interval

$[t, t + dt]$ is given by Eq. 3.

$$dp = \lambda_R dt$$

In this subsection, additional variations of accident frequencies with time are not considered, to avoid complicating the analysis. In addition, the probability of an accident in an interval of length dt in the complement of the region (the rest of the world), $\widetilde{R}$, is given by

$$dp = \widetilde{\lambda}_R dt$$

Assuming a flat prior distribution for $\widetilde{\lambda}_R$, a posterior probability distribution for this parameter is now constructed. Assuming that $n_{\tilde{R}}$ accidents have been observed in the rest of the world in a total operating time $T_{\tilde{R}}$, and using techniques presented previously, the posterior probability distribution for $\widetilde{\lambda}_R$, is given by

$$\mathcal{P}(\widetilde{\lambda}_R|\widetilde{n}_R) = \frac{1}{\Gamma(\widetilde{n}_R + 1)} \widetilde{T}_R (\widetilde{\lambda}_R \widetilde{T}_R)^{\widetilde{n}_R} e^{-\widetilde{\lambda}_R \widetilde{T}_R}.$$

Introducing the assumption that the prior distribution for λ_R is the same as the posterior distribution for $\widetilde{\lambda}_R$, except for a constant of proportionality κ between them, the factor κ indicates our prior belief that nuclear safety in region R is better or worse than that in other parts of the world.

For example, one could obtain an estimate for κ using a record of less severe incidents (not necessarily core-damage or large-release) that have occurred in R and $\widetilde{R}$

$$\kappa = \frac{I_R}{\widetilde{I}_R},$$

where I_R, $\widetilde{I}_R$ are the number of incidents recorded in R and $\widetilde{R}$ respectively. It is, of course, necessary to have a precise criterion to count the incidents above, and one example is given below. Since incidents of lower severity are fairly frequent, a purely frequentist analysis is sufficient to obtain an estimate for κ.

After fixing κ using this technique, or some other, the prior distribution for λ_R is taken as

$$\mathcal{P}(\lambda_R) = \frac{1}{\kappa} \mathcal{P}(\widetilde{\lambda}_R|\widetilde{n}_R)\Big|_{\widetilde{\lambda}_R = \frac{\lambda_R}{\kappa}} = \frac{1}{\kappa^{\widetilde{n}_R+1}\Gamma(\widetilde{n}_R + 1)} \widetilde{T}_R (\lambda_R \widetilde{T}_R)^{\widetilde{n}_R} e^{-\frac{\lambda_R \widetilde{T}_R}{\kappa}} \tag{12}$$

Now, denoting the number of accidents observed inside region R by n_R, in a total operating time T_R, a posterior probability distribution for λ_R can be constructed using the same techniques and the prior in (12). This leads to

$$\mathcal{P}(\lambda_R|n_R) = \frac{T_R + \frac{\widetilde{T}_R}{\kappa}}{\Gamma(n_R + \widetilde{n}_R + 1)} \left(\left(T_R + \frac{\widetilde{T}_R}{\kappa} \right) \lambda_R \right)^{\widetilde{n}_R + n_R} e^{-\lambda_R (T_R + \frac{\widetilde{T}_R}{\kappa})}$$

The probability for the hypothesis that $\lambda_R < \lambda_0$ is given by

$$\mathcal{C}_R(\lambda_0, n_R) = \frac{\Gamma(1 + n_R + \widetilde{n}_R, (T_R + \frac{\widetilde{T}_R}{\kappa})\lambda_0)}{\Gamma(1 + n_R + \widetilde{n}_R)}.$$

In the end, the model yields a simple result. The rule is that to account for variations in safety across regions, one simply counts the total accidents $\widetilde{n}_R + n_R$ as having occurred in a time $T_R + \frac{\widetilde{T}_R}{\kappa}$. In the situation where $\kappa = 1$, the results of the simple model are obtained again since $T_R + \widetilde{T}_R$ just becomes the total operating experience T_{obs}.

Recalling the expansion of the incomplete gamma function shown in equatio (8), and for small $\lambda_0(T_R + \frac{\widetilde{T}_R}{\kappa})$ the expressions above are well approximated by

$$\mathcal{C}_R(\lambda_0, n_R) \approx \frac{\left(\lambda_0(T_R + \frac{\widetilde{T}_R}{\kappa})\right)^{n_R+\widetilde{n}_R+1}}{\Gamma(n_R + \widetilde{n}_R + 2)}$$

If λ_0 is set to one of the values given in Table 1, the expression above necessarily evaluates to an extremely small number.

As an example, France is taken as the region under consideration. At the end of 2012, France had accumulated about $T_R = 1874$ reactor-years of operating experience, without a single large-release event.[18] In the same period the rest of the world had accumulated $\widetilde{T}_R = T_{\text{obs}} - T_R = 13{,}373$ reactor-years of operating experience. Furthermore, κ is set to 0.5, which suggests a prior belief that French reactors are twice as safe as the world-average (this value is just taken to illustrate this example, and do not suggest that this is really the case).

With these parameters, given that there have been five large release events in the rest of the world, and taking the PRA frequency estimate for the French EPR reactor given above, the following value for $\mathcal{C}_R$ is obtained

$$\mathcal{C}_R(7.69 \times 10^{-8}, 5) = 1.6 \times 10^{-19}.$$

Therefore, given the existing empirical experience in the rest of the world, even with an assumption that French reactors are considerably safer, the probability that the EPR reactor genuinely has a true frequency of accidents as small as the predictions of PRA is absurdly low.

It is clear that the models in this section can be extended further to account for more detailed variations. However, the analysis of this section shows that the severe contradiction between the results of PRA and empirical data cannot be resolved in any such manner.

On the Indian experience

This section discusses the Indian experience. The correct procedure to analyze this case would be to use the rest of the world's experience as a prior distribution as was done in the previous section. Needless to say, with any reasonable choice of κ, as defined there, it can be concluded that for India, $\epsilon \ll 1$. However, in this short subsection, a separate approach is taken to make an elementary statistical point. Even if one assumes that the existing record of nuclear accidents has absolutely no bearing on the Indian situation, India's operating experience of $T_{\text{ind}} = 394$ reactor years

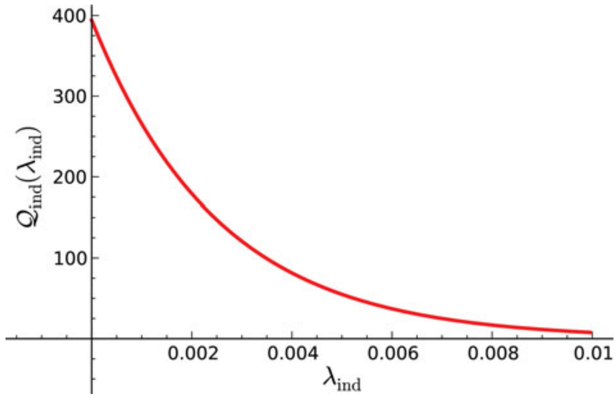

Figure 5. Posterior probability distribution for λ_{ind}.

is too low to provide any statistical confidence in the safety of the Indian nuclear programme.[28]

The significance of this observation pertains to common claims made by the Indian nuclear establishment that Indian reactors are safer since India has not witnessed a major accident in this time period. The point of this section is to point out that such claims are statistically fallacious. While the discussion is phrased in India's context, it is worth noting that several countries have accumulated similar levels of operating experience. To take a few examples, by the end of 2012, Belgium had accumulated 254 reactor-years, China had accumulated 141 reactor-years, Canada had accumulated 634 reactor-years, the republic of Korea had accumulated 404. Although none of these countries have seen a major accident, the conclusions of this section apply to all of them: their operating experience is insufficient to suggest that safety levels in these countries are significantly different from the world-average.

In mathematical terms, in this section the objective is somewhat different from the rest of the article. Here, the goal is not to estimate ϵ (from Eq. 1) and show that it is very small, but rather to show that even if one discards the reasonable prior assumptions made in the discussion on regional variations, there is no reason to suppose that $\epsilon \sim 1$ for these countries. As explained above, the numerical figures below are specific to India, but similar results can be computed for different countries using the same calculations.

There have been several minor but no major accidents in India's operating history. Starting with a flat prior distribution for the mean frequency of accidents at Indian reactors, denoted by λ_{ind} to distinguish it from the global frequency, the posterior distribution for λ_{ind} is given by

$$Q_{ind}(\lambda_{ind}) = T_{ind} e^{-\lambda_{ind} T_{ind}},$$

Note that this result can be obtained by setting $n_{obs} \to 0$ and $T_{obs} \to T_{ind}$ in Eq. 6. This curve is plotted in Figure 5.

This figure tapers off quite gently, so India's current operating experience cannot tell much about the frequency of accidents in India, especially if the true frequency is around ν_{obs}^{cd} in (2).

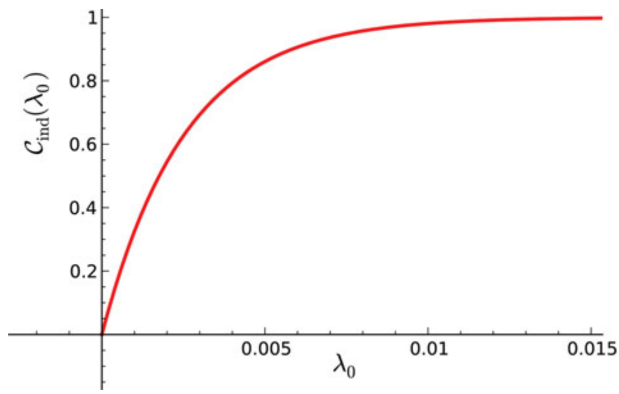

Figure 6. Probability for the hypothesis $\lambda_{\text{ind}} < \lambda_0$.

To make this sharper, it is useful to look at the probability for the hypothesis that $\lambda_{\text{ind}} < \lambda_0$. This probability is given by the function

$$\mathcal{C}_{\text{ind}}(\lambda_0) = \int_0^{\lambda_0} \mathcal{Q}_{\text{ind}}(\lambda)\, d\lambda = 1 - e^{-T_{\text{ind}}\lambda_0}.$$

This curve is shown in Figure 6.

Some relevant numerical figures are

$$\mathcal{C}_{\text{ind}}(2.67 \times 10^{-4}) \approx 0.10,$$
$$\mathcal{C}_{\text{ind}}(1.0 \times 10^{-3}) \approx 0.33.$$

In words, this indicates that it is only with a confidence of about 10% that one can state that the true frequency of accidents is smaller than once in 3,700 reactor-years. And it is only with a confidence of about 33% that one can conclude that the true frequency of accidents in India is smaller than once per thousand reactor-years.

These results are very simple to understand intuitively. They simply reflect the fact that even if the expected frequency of accidents in India is once every thousand reactor-years, there is still an excellent chance that one could get through 394 reactor-years without an accident. Conversely, the absence of an accident in this time period is not particularly informative.

It is quite common to find claims like the one made by the Indian Nuclear Society its advertisement for a 2012 conference. "Having achieved safe and reliable operation of about 360 reactor-years . . . the Indian nuclear programme has demonstrated a high level of maturity. The safety track record of Indian nuclear power plants has been impeccable."[29] These claims are accepted at the highest levels of government. In 2011, the Prime Minister stated that "The safety track record of our nuclear power plants over the past 335 reactor-years of operation has been impeccable."[30] The analysis presented in this article shows that it is erroneous to draw these complacent conclusions. India's operating experience is too limited to provide statistically reliable long term estimates about the efficacy or otherwise of safety practices in the nuclear sector.

In passing, it can be pointed out that, within the nuclear industry, it appears to be a rather common statistical error to extrapolate from limited experience to rather strong claims of safety. For example, the Rasmussen report also stated that "It is significant that in some 200 reactor-years of commercial operation of reactors of the type considered in the report there have been no fuel melting accidents."[11] However, this experience had absolutely no significance for the conclusion drawn by the report, which was that the core-damage frequency of reactors was once in 20,000 years.

The debate on nuclear liability

It is clear that the results of PRA are untenable in the light of empirical data. In this section, using the debate on nuclear liability, evidence is provided that the industry's actions (as opposed to its public statements) suggest that it has independently reached this conclusion.

Soon after the nuclear deal, multinational nuclear suppliers lobbied the government to pass a law that would indemnify them in the event of an accident. Under this pressure, the government passed a liability law in 2010 that was almost identical to the annex of a U.S. sponsored international convention on the subject called the Convention on Supplementary Compensation.[31,32] However, as a result of various pulls-and-pushes in the legislative process, nuclear suppliers are largely protected but not completely indemnified by the Indian law. While victims cannot sue the supplier, the law allows the NPCIL, which has the primary responsibility of compensating victims, to recover some of this compensation from the supplier for a disaster caused by substandard equipment. The refusal of suppliers to accept even this marginal liability has presented a significant obstacle for contracts for new reactors (see for example, a recent statement by the CEO of General Electric).[33]

It is significant that although one of primary advertised benefits of the Indo–U.S. nuclear deal was that India would be able to purchase light water reactors from new suppliers, the deadlock on liability has prevented this entirely. The first new contract for reactors after the nuclear deal was signed in 2014 with the Russian public sector company Rosatom. However, this will just extend the Kudankulam nuclear complex which was covered by an inter-governmental agreement even before the nuclear deal. Negotiations with suppliers from markets that have opened up after the nuclear deal, including companies from France and the United States, have so far failed to yield a single new contract due to the conflict on liability.

However, this leads directly to the following question: if the chance of a nuclear accident is indeed as remote as the industry claims, then why are nuclear reactor manufacturers unwilling to accept liability for an accident?

One of the ostensible reasons given by suppliers is that forcing them to accept liability would cause the cost of power to go up.[34]

To examine the veracity of this claim, consider the expected cost of insurance for suppliers if the results of PRA are taken at face value by the industry and actuaries.[35]

The Indian liability law currently caps the total available compensation for victims at "the rupee equivalent of three hundred million Special Drawing Rights"(SDRs).[36] This includes compensation from the operator, and also the central government, and victims are not legally entitled to any further compensation. Here the worst case scenario, for the supplier, is considered. This is when it becomes liable for this entire amount.

Although the rupee to SDR exchange rate fluctuates, this maximum liability is approximately $\ell_{\text{cap}} = \text{Rs. } 2,500 \text{ crores} = \text{Rs. } 2.5 \times 10^{10}$ in rupee terms.

Taking, for example, the claimed frequency of core-damage accidents at the Kudankulam reactors, which is denoted by $\mu_{\text{kk}} = 10^{-7}(\text{ry})^{-1}$, then a simple order of magnitude estimate for the cost of insurance for this amount is

$$i_p = \mu_{\text{kk}} \times \ell_{\text{cap}} = \text{Rs. } 2,500.$$

A reactor with a capacity of 1000 MW, operating at a 80% load-factor, should produce $E = 0.8 \times 10^6 \times 365 \times 24$ kWh $\approx 7 \times 10^9$ kWh of electricity each year. So, the cost of insurance above should lead to an increase in the cost of electricity by

$$\delta p = \frac{i_p}{E} = 3.6 \times 10^{-7} \text{ Rs./kWh!}$$

This absurdly small number indicates that something is amiss with the industry's claim that liability will lead to price increases.

In fact two factors of about about 10^3 and 10^4, which are missing in the calculation above, are required to make sense of the industry's reluctance. The first is that nuclear accidents could lead to damage that is a thousand times more than the cap on liability. For example, some estimates of the economic damage at Fukushima are as high as $\ell_{\text{real}} \approx$ USD 200 billion $\approx$ Rs. 12 lakh crores.[37]

In principle, a future Indian government could ignore the liability cap and insist on recovering larger costs from the supplier. Even this would not lead to a prohibitive cost of insurance if reactors were genuinely as safe as manufacturers' claim. The other crucial missing factor comes from our result above: accidents affecting the public are likely to happen at a rate that is closer to $\nu_{\text{obs}}^{\text{lr}}$. Extending this simple linear model with these realistic estimates of damage and risk leads to the following cost of insurance per unit energy produced

$$\delta p_{\text{true}} = \frac{\nu_{\text{obs}}^{\text{lr}} \times \ell_{\text{real}}}{E} = 0.56 \text{ Rs./Kwh.}$$

This is now a significant fraction (roughly 10%) of the cost of electricity. However, at this point, corrections to this linear model for the insurance premium become significant. For example, since the total amount involved, ℓ_{real}, is very high, and the expected rate $\nu_{\text{obs}}^{\text{lr}}$ is non-negligible, financial institutions would evidently be unwilling to underwrite this risk without additional incentives in the form of a significantly higher cost of insurance. This helps explain why suppliers insist on legislative indemnity, rather than simply arranging for the appropriate financial cover.

What the debate on liability shows is that the nuclear industry—both in the private and the public sector—has itself taken note of the empirical rates of accidents, and it is unwilling to take the predictions of its PRAs seriously when its economic interests are at stake.

Conclusions

This article, by means of some simple Bayesian calculations, reaches the following conclusion: the historical record contradicts the predictions of probabilistic risk assessment and suggests a significantly higher frequency of nuclear accidents. The contradiction between these predictions and data can be quantified in terms of a probability for the hypothesis that true frequencies of accidents are as small as those predicted. This probability can be quantified in various models, and is found to be extraordinarily small in a model-independent fashion, and independent of changes in our detailed assumptions. In particular, it showed how, even in models where safety standards improve exponentially with operating experience or where reactors in a given region are assumed to be considerably safer than reactors elsewhere, the conclusions above hold at very high confidence levels.

Second, by analyzing in depth the case of India (strictly speaking a subset of the study of worldwide accident frequencies above), it showed that India's current experience with nuclear reactor operation is far from sufficient to draw any strong conclusions about future reactor safety. The same conclusions hold for other countries with similar levels of operating experience, such as Canada, China, Belgium and Korea.

Therefore it is clear that the methodology and practice of PRA needs to be revised significantly.

In fact, as we showed in the discussion on accident liability, it is clear that the nuclear industry already recognizes that the results of PRA are numerically unreliable. Nevertheless, within the technical community, the use of PRA is sometimes justified as a useful tool for safety analysis. For example, it has been explained elsewhere that even though PRA is "not thought to represent the true risks" it remains useful as a "platform for technical exchanges on safety matters between regulators and the industry."[9] Indeed, the Rasmussen report started by explaining that "the objective of the study was to make a realistic estimate of [the] risks" that would be "involved in potential accidents."[11]

Apart from the theoretical problems mentioned in the introduction, it appears likely that the nuclear industry benefits from the disingenuous suggestion that it can, in fact, accurately predict the frequency of accidents. Although insiders recognize that this is not the case, it is clear that the detailed computer simulations that support a supposedly scientific calculation of low accident frequencies computed to several decimal places are useful in public debates. While there have been other critiques of the mismatch between the results of PRA and empirical data, this study is significant because it emphasizes the extraordinarily high level of confidence with which it is possible to rule out the results of PRA.

Indeed, this would hardly be the only situation in which the nuclear industry has attempted to use the authority of science to dismiss safety concerns. For example, in an attempt to dismiss the history of Chernobyl, the World Nuclear Association declared in January 2011 that "In the light of better understanding of the *physics and chemistry* of material in a reactor core . . . it became evident that even a severe core melt coupled with breach of containment could not in fact create a major radiological disaster from any Western reactor design" (emphasis added).[38] After attempting to initially defend this claim for a few days after Fukushima by claiming that "clearly there was no major release from the reactors" and only from the "fuel pools,"[39] the Association had to reluctantly concede that its claim, seemingly based on rigorous material science, "did not apply to all" Western reactor designs.[40]

It is interesting to note that a similar dynamic operates in other industries as well. For example, in the aviation industry (which inspired the use of PRA for nuclear reactors), as part of the certification process for its new 787 "Dreamliner" aircraft, Boeing estimated that its lithium-ion batteries would vent smoke "once in every 10 million flight hours." In fact this event occurred twice in 52,000 flight hours leading to the grounding of the entire fleet for inspection.[41]

To explore the implications of these conclusions, it is worth returning to the case of India where the government is planning a large nuclear expansion.

It has announced plans to commence construction on eight heavy water reactors, with a capacity of 5600 MW in the "12th Plan" period (2012–17), and complete work on a separate 2800 MW of installed capacity. In addition, it is also planning to import eight reactors with a total capacity of 10,500 MW.[42]

Every reactor site has seen vigorous local protest movements that have raised issues of land and livelihood as well as questions about nuclear safety.

Therefore it is imperative to have a frank conversation on nuclear safety, involving not just the technical community but a far broader cross-section of society. The results presented in this article show that such a debate should start with the acceptance that the ambitious claims about nuclear safety made on the basis of probabilistic risk assessment have been conclusively falsified by the empirical data.

Acknowledgments

This article is based on a detailed exchange of letters with Mr. Nalinish Nagaich of the NPCIL. I am grateful to M. V. Ramana for comments on a draft of this manuscript.

References

1. Press Trust of India, "AEC Chief puts Odds of N-Plant Accidents at 1-in-infinity," *The Hindu*, 10 November 2011. http://www.thehindu.com/news/national/aec-chief-puts-odds-of-nplant-accidents-at-1ininfinity/article2615375.ece.
2. S. Agrawal, A. Chauhan, and A. Mishra, "The VVERs at KudanKulam," *Nuclear Engineering and Design*, 236(7), 2006: 812–835. http://dx.doi.org/10.1016/j.nucengdes.2005.09.030.

3. Areva and Electricite de France, *Pre Construction Safety Report of the UK EPR*, ch. 15.7: PSA Discussion and Conclusions. U.K. Health and Safety Executive, 2012: 14. http://www.epr-reactor.co.uk/ssmod/liblocal/docs/PCSR/Chapter%2015%20-%20Probabilistic%20Safety%20Analysis/Sub-Chapter%2015.7%20-%20PSA%20Discussion%20and%20Conclusions.pdf.
4. Westinghouse, *Pre Construction Safety Report of the AP 1000*. UK Health and Safety Executive, 2009. https://www.ukap1000application.com/PDFDocs/Safety/UKP-GW-GL-732%20Rev%201.pdf.
5,. Nuclear Regulatory Commission, "New Light-Water Reactors Sub-arena: Goals.". http://www.nrc.gov/about-nrc/regulatory/risk-informed/rpp/reactor-safety-new.html.
6. Dennis Bley, Mary Drouin, Tom King, Jeffrey LaChance, John Lehner, Bruce Mrowca, and Vinod Mubayi, *Feasibility Study for a Risk-Informed and Performance-Based Regulatory Structure for Future Plant Licensing*. U.S. Nuclear Regulatory Commission, December 2007. http://pbadupws.nrc.gov/docs/ML0804/ML080440215.pdf.
7. Working Group on Risk Assessment, *Probabilistic Risk Criteria and Safety Goals*. Nuclear Energy Agency (OECD), 17 December 2009. https://www.oecd-nea.org/nsd/docs/2009/csni-r2009-16.pdf.
8. Atomic Energy Regulatory Board, *Consenting Process for Nuclear Power Plants and Research Reactors*, 2007. http://www.aerb.gov.in/AERBPortal/pages/English/t/publications/CODESGUIDES/SG-G-1.pdf.
9,. D. Sornette, T. Maillart, and W. Kröger, "Exploring the Limits of Safety Analysis in Complex Technological Systems," *International Journal of Disaster Risk Reduction*, 6, (2013): 59–66.
10. M. Ha-Duong and V. Journé, "Calculating Nuclear Accident Probabilities from Empirical Frequencies," *Environment Systems and Decisions*, 34(2), 2014: 249–258.
11. Norman Rasmussen et al., *Reactor Safety Study: an Assessment of Accident Risk in U.S. Commercial Nuclear Power Plants*. Nuclear Regulatory Commission, 1975. http://www.osti.gov/scitech/servlets/purl/7134131.
12. Kenneth Kok, ed., *Nuclear Engineering Handbook*. CRC Press, 2009.
13. W. Keller and M. Modarres, "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, (2005): 271–285. http://dx.doi.org/10.1016/j.ress.2004.08.022.
14. World Nuclear News, "Earthquake not a Factor in Fukushima Accident," 2 December 2011. http://www.world-nuclear-news.org/RS_Earthquake_not_a_factor_in_Fukushima_accident_0212111.html.
15. François Diaz Maurin, "*The Viability and Desirability of Alternative Energy Sources: Exploring the Controversy over Nuclear Power*." Universitat Autónoma de Barcelona, 2013. PhD dissertation. https://www.educacion.gob.es/teseo/imprimirFicheroTesis.do?fichero=44006.
16. Areva and Electricite de France, *Pre Construction Safety Report of the UK EPR*, ch. 15.4: Level 2 PSA. 222. U.K. Health and Safety Executive, 2012. http://www.epr-reactor.co.uk/ssmod/liblocal/docs/PCSR/Chapter%2015%20-%20Probabilistic%20Safety%20Analysis/Sub-Chapter%2015.4%20-%20Level% 202% 20PSA.pdf.
17. World Nuclear Association, "Advanced Nuclear Power Reactors," April 2014. http://www.world-nuclear.org/info/Nuclear-Fuel-Cycle/Power-Reactors/Advanced-Nuclear-Power-Reactors/.
18. International Atomic Energy Agency, *Nuclear Power Reactors in the World*, 2013. http://www-pub.iaea.org/MTCD/Publications/PDF/rds2-33_web.pdf.
19. Thomas B. Cochran and Matthew G. McKinzie, "Global Implications of the Fukushima Disaster for Nuclear Power," in *World Federation of Scientists International Seminar on Planetary Emergencies*, August 2011. http://docs.nrdc.org/nuclear/files/nuc_11102801a.pdf.

20. H.W. Bertini and Staff of the Nuclear safety Information Center, *Descriptions of Selected Accidents that Have Occurred at Nuclear Reactor Facilities*. Oak Ridge National Laborator, 1980. p. 35. http://www.nucleartourist.com/references/Descriptions%20of%20Selected%20Nuclear%20Reactor%20Accidents.pdf.
21. John G Kemeny et al., *The Need for Change, the Legacy of TMI: Report of the President's Commission on the Accident at Three Mile Island*. US Government Printing Office, 1979. http://www.threemileisland.org/downloads/188.pdf.
22. International Atomic Energy Agency, *International Nuclear and Radiological Event Scale*, 2008. http://www-pub.iaea.org/MTCD/publications/PDF/INES2009_web.pdf.
23. International Nuclear Safety Advisory Group, *The Chernobyl Accident: Updating of INSAG-1*. No. 75 in Safety Series, International Atomic Energy Agency, 1992. http://www-pub.iaea.org/MTCD/publications/PDF/Pub913e_web.pdf.
24. World Nuclear Association, *Safety of Nuclear Power Reactors*, April 2014. http://www.world-nuclear.org/info/Safety-and-Security/Safety-of-Plants / Safety - of-Nuclear-Power-Reactors/.
25. World Nuclear Association, *Decommissioning Nuclear Facilities*, March 2014. see Table "Reactors closed following damage in an accident or serious incident." http://www.world-nuclear.org/info/nuclear-fuel-cycle/nuclear-wastes/decommissioning-nuclear-facilities/.
26. WISE Amsterdam and WISE Stockholm, "Near Core Meltdown Covered up in GDR," *World Information Service on Energy*, 9 February 1990. http://www.wiseinternational.org/node/133.
27. Kiyoshi Kurokawa et al., *Fukushima Nuclear Accident Independent Investigation Commission*. National Diet of Japan, 2012. https://www.nirs.org/fukushima/naiic_report.pdf.
28. Nuclear Power Corporation of India Limited. "Safety Performance." http://www.npcil.nic.in/main/aboutus.aspx.
29. Indian Nuclear Society. "Safety in Design, Construction and Operation of Nuclear Power Plants." http://indiannuclearsociety.in/conference/.
30. Manmohan Singh, "Remarks by the Prime Minister on the Occasion of the Conferment of the Department of Atomic Energy's Lifetime Achievement Awards 2009," 29 March 2011. http://pib.nic.in/newsite/PrintRelease.aspx?relid=71343.
31. Suvrat Raju and M. V. Ramana, "The Other Side of Nuclear Liability," *Economic and Political Weekly*, vol. XLV, 17 April 2010. http://www.princeton.edu/ ramana/NuclearLiability-Vol45-No16-17April2010.pdf.
32. Suvrat Raju and M. V. Ramana, "Strange Love," *OPEN Magazine*, May 2011. http://www.openthemagazine.com/article/nation/strange-love.
33. Sunil Jagtiani and N. C. Bipindra, "GE Won't Build India Nuclear Plants on Liability Risk, Says Immelt," *Mint*, 21 September 2015. http://www.livemint.com/Industry/5RVZfc57Fg4UwKtiT1wSJI/ GE-wont-build-India-nuclear-plants-on-liability-risk-says.html.
34. Express News Service, "More Liabilities Will Push up Kudankulam Cost: Russia," *Indian Express*, 15 October 2012. http://archive.indianexpress.com/news/more-liabilities-will-push-up-kudankulam-cost-russia/1016840/.
35. M. V. Ramana and Suvrat Raju, "Nuclear Safety Before Vendor Interests," *Hindu*, 30 October 2012. http://www.thehindu.com/opinion/op-ed/nuclear-safety-before-vendor-interests/article4044438.ece.
36. Ministry of Law and Justice, "The Civil Liability for Nuclear Damage Act 2010," *The Gazette of India (extraordinary)*, 22 September 2010. http://www.prsindia.org/uploads/media/Nuclear%20Rules/The%20Civil%20Liability%20for%20Nuclear%20Damage%20Act.pdf.
37. JCER Economic Research Department, "Impact to Last Decade or More if Existing Nuclear Plants Shut Down," Tech. Rep., Japan Center for Economic Research, 25 April 2011. http://www.jcer.or.jp/eng/research/pdf/pe(iwata20110425)e.pdf.

38. World Nuclear Association. "Safety of Nuclear Power Reactors," [online]. January 2011. https://web.archive.org/web/20110220001120/http://www.world-nuclear.org/info/inf06.html. Archived at the Internet Archive "Wayback Machine."
39. World Nuclear Association. "Safety of Nuclear Power Reactors," [online]. March 2011 [accessed 12 October 2015]. Available from: https://web.archive.org/web/20110407010308/http://www.world-nuclear.org/info/inf06.html. Archived at the Internet Archive "Wayback Machine."
40. World Nuclear Association. "Safety of Nuclear Power Reactors," [online]. June 2011. https://web.archive.org/web/20110622094525/http://www.world-nuclear.org/info/inf06.html. Archived at the Internet Archive "Wayback Machine."
41. Office of Aviation Safety, *Interim Factual Report on case DCA13IA037 involving a Japan Airlines 787-8*. National Transportation Safety Board, 7 March 2013. http://www.ntsb.gov/investigations/2013/boeing_787/interim_report_B787_3-7-13.pdf.
42. V. Narayanasamy, "Setting up of Nuclear Plants in Eleventh/Twelfth Plans: Answer to Unstarred Question no. 316 in the Rajya Sabha," 15 March 2012. http://dae.nic.in/writereaddata/rsus316.pdf.