

COWLES FOUNDATION FOR RESEARCH IN ECONOMICS

AT YALE UNIVERSITY

Box 2125, Yale Station  
New Haven, Connecticut 06520

COWLES FOUNDATION DISCUSSION PAPER NO. 532

Note: Cowles Foundation Discussion Papers are preliminary materials circulated to stimulate discussion and critical comment. Requests for single copies of a Paper will be filled by the Cowles Foundation within the limits of the supply. References in publications to Discussion Papers (other than mere acknowledgment by a writer that he has access to such unpublished material) should be cleared with the author to protect the tentative character of these papers.

HOW SHOULD WE REVISE OUR BELIEFS ABOUT  
NUCLEAR POWER SAFETY AFTER THREE MILE ISLAND?

William D. Nordhaus

July 24, 1979

HOW SHOULD WE REVISE OUR BELIEFS ABOUT  
NUCLEAR POWER SAFETY AFTER THREE MILE ISLAND?

Summary

In light of the accident at Three Mile Island, the paper presents a preliminary analysis of the compatibility between the analytical work of the Reactor Safety Study (Rasmussen Report) and actuarial experience. The technique is a "macroanalytic" approach rather than the "microanalytic" method of the Reactor Safety Study.

The first question asked is how likely it is that an accident as severe as that at Three Mile Island would occur if the Reactor Safety Study is correct. Using the most likely estimates, it is concluded that the chances are 1 in 80 that such an accident would occur this soon, but given uncertainties about parameters it might range from 1 in 17 to 1 in 625.

The second question addressed is how we should revise our estimates of the safety of nuclear power given the Three Mile Island experience. Using the technique of maximum likelihood, our best guess estimate of the risk of accidents causing at least one fatality rises from the Reactor Safety Study's 32 per million reactor years to about 2000 per million reactor years.

## 1. Introduction

In the well-known Reactor Safety Study<sup>1</sup> (or Rasmussen Report), an attempt was made to estimate the probability of serious accidents in light water reactors. The basic methodology was to identify alternative routes by which accidents occur and then to assess the probability of each. \*

Figure 1, from the introduction of the final report, illustrates the frequency-severity function that the report derived.

The Reactor Safety Study has been heavily criticized on several grounds. Some critics questioned the methodology; others the detailed technical judgments; others the author's objectivity. In my own view, it represented the most careful attempt to quantify the risks, but given the paucity of data the most careful report could be far off the mark. The fundamental weakness is simply that we have insufficient actual experience to judge whether the views were correct. And the ultimate test would come when sufficient time and experience proved the analysis correct or deficient.

---

\*I would like to thank John Hartigan and Tjalling Koopmans for helpful comments and Edward Buffie for research assistance.

<sup>1</sup>U. S. Nuclear Regulatory Commission, Reactor Safety Study: An Assessment of Accident Risks in U. S. Commercial Nuclear Power Plants, WASH-1400 (NUREG-75/014), October 1975.




FIGURE 1-1 Frequency of Fatalities due to Man-Caused Events

- Notes
1. Fatalities due to auto accidents are not shown because data are not available. Auto accidents cause about 50,000 fatalities per year.
  2. Approximate uncertainties for nuclear events are estimated to be represented by factors of 1/4 and 4 on consequence magnitudes and by factors of 1/5 and 5 on probabilities.
  3. For natural and man caused occurrences the uncertainty in probability of largest recorded consequence magnitude is estimated to be represented by factors of 1/20 and 5. Smaller magnitudes have less uncertainty.




FIGURE 1-2 Frequency of Fatalities due to Natural Events

- Notes
1. For natural and man caused occurrences the uncertainty in probability of largest recorded consequence magnitude is estimated to be represented by factors of 1/20 and 5. Smaller magnitudes have less uncertainty.
  2. Approximate uncertainties for nuclear events are estimated to be represented by factors of 1/4 and 4 on consequence magnitudes and by factors of 1/5 and 5 on probabilities.

Figure 1. Estimates of risk from nuclear power and other events from Reactor Safety Study

Source: Reactor Safety Study, p. 2.

The present note attempts to show how it is possible to integrate actual experience with the analytical estimates of reactor safety. We will ask two simple questions:

1. How likely is it that an accident like that at Three Mile Island would occur if the Rasmussen estimates are correct?
2. How should we revise our estimates of the safety of nuclear power given the Three Mile Island accident?

It is important to note that the object of this discussion is to examine the overall plausibility of the risk estimates from the Reactor Safety Study. It can be called a "macrostatistical" approach as opposed to the Reactor Safety Study's "microstatistical" approach. In the latter, all the important faults that could be identified were enumerated and (crudely speaking) their probabilities were added up. The overall risk estimate is thus based on the estimates of risks of individual components. A macrostatistical approach looks at the overall failure rate without enquiring into the components of the failure.

This methodological distinction is important because the ultimate objective of the Reactor Safety Study was to judge the overall risk. If one is attempting to choose between alternative energy technologies, it matters little if the study proves defective because the probabilities were incorrect, if a particular event chain were ignored, or if a certain form of human failure were ignored. The charge of the study was to assess relevant risk, not the risk of irrelevant events.

## 2. Consistency of Rasmussen and Three Mile Island

The Reactor Safety Study has calculated the frequency of accidents of given severity (shown in Figure 1). If accidents are caused by component or systems failure, the chance of an accident occurring is independent of what happens at other times or at other reactors. In this case, the waiting time until an accident occurs has an exponential distribution<sup>2</sup>. More precisely, if a given reactor has operated for  $T$  years, the probability that an accident has not yet occurred will be:

$$(1) \quad p(T) = e^{-\alpha T}, \quad \alpha > 0.$$

More generally,  $T$  can be interpreted as the number of reactor-years of operation to date. In this distribution, the expected number of reactor-years between accidents is equal to  $1/\alpha$ .

The distribution for the frequency of accidents can be broadened to include accidents of different "severity" levels. Thus, let  $D$  be the number of total deaths ultimately caused by an accident. Then we have a family of distributions relating frequency of accidents to severity, where the parameter  $\alpha$  depends on the severity level that is examined.

---

<sup>2</sup>The assumption that the "macro" accidents have an exponential distribution is not investigated by the Reactor Safety Study. The distribution of failures is assumed to be exponential for "nonreparable components," while "reparable components" and combined failures follow standard reliability theory. In the latter case, the exponential distribution will, however, be a good approximation as long as the probability of failure is small. Finally, note that if all individual faults or components are exponentially and independently distributed, the overall "macro" accident will have an exponential distribution.

According to the Reactor Safety Study, the frequencies of different accidents can be represented by the following:<sup>3</sup>

Severity (D) (fatalities)	Frequency (per million reactor years, MRY) <u>optimistic</u>	<u>most likely</u>	<u>pessimistic</u>
At least 1	6.3	32.	160.
10	4.0	20.	100.
100	0.50	2.5	13.
1000	0.006	0.032	0.16

Given such a distribution, how likely was the Three Mile Island accident? According to HEW Secretary Califano and an U. S. Government working group, the ultimate number of deaths is expected to be 1, but might range as high as 10.<sup>4</sup> Presumably, there will be some further refinements of these estimates in the course of the Kemeny Commission.

Given these data, we ask what is the probability that the Three Mile Island accident "should" have occurred. More precisely, assume (a) that the Califano estimate is correct, (b) that the Reactor Safety Study most likely estimate is correct, and (c) that the underlying distribution is exponential. This implies:<sup>5</sup>

The chance that an accident as severe as Three Mile Island would occur if the Reactor Safety Study estimates are correct is approximately 1 in 80.

<sup>3</sup>The estimates include both early and latent cancer fatalities from Reactor Safety Study, pp. 88 and 90.

<sup>4</sup>New York Times, May 13, 1979.

<sup>5</sup>The calculation behind this statement is as follows: For  $D \geq 1$ , the Rasmussen estimate is that  $\alpha$  (probability of an accident per reactor-year) =  $32. \times 10^{-6}$ . There have been approximately 400 reactor-years of experience in the United States. Therefore the probability of an accident in the first 400 reactor-years is

$$p = 1 - \exp[-400 \times 32 \times 10^{-6}] = 0.013$$

By the same technique, we can take into account the fact that there is considerable uncertainty about the consequences of the Three Mile Island accident and that the Reactor Safety Study gave pessimistic and optimistic probability estimates. The following table gives the probabilities that the Three Mile Island accident would occur given alternative estimates of the probabilities or of the consequences.

Table 1. Chance that Three Mile Island Accident would Occur if Rasmussen Estimates are Correct

	Alternative Consequences of Three Mile Island (fatalities)	
	1.0	10.0
	Alternative Probabilities of Events in Reactor Safety Study	
Optimistic	1 in 417	1 in 625
Most Likely	1 in 80	1 in 125
Pessimistic	1 in 17	1 in 25

The results of this calculation are straightforward. According to the Reactor Safety Study, given the existing number of reactors in the United States, an event as severe as Three Mile Island should have occurred once every few hundred years rather than once every ten years. The "pessimistic" probability estimates are just marginally better.

Put differently, the chance that such an accident would occur is approximately equivalent to flipping a fair coin heads six or seven times in a row. Faced with such a string of heads, one would surely want to peek



on the other side to see whether both sides are heads! The only plausible conclusion is that the Rasmussen risk estimates cannot be accepted as realistic for public policy purposes.

### 3. Revised Estimates of Accidents

Clearly, the risks of nuclear power (and of believing the experts) will require a thorough review over the coming years. While this detailed "microanalysis" progresses, we might ask, is there a short-cut or "macro-analytical" technique for integrating the prior knowledge we had (the Reactor Safety Study) with what we have observed (the actual accident record).

The problem we face is illustrated in Figure 2. The solid line shows the risk estimate from the Reactor Safety Study. The box shows the "best guess" of the risk of accidents causing at least one fatality, based on actuarial experience. Our dilemma is that if we place a great deal of confidence in theory, our assessments have not changed very much. On the other hand, if we distrust theory and believe only history, then our assessment will change enormously.

In what follows we ask, given both theory and history, what is the most likely assessment of the risk of severe nuclear accidents? The technique used to answer this question is the well-worn procedure called "maximum likelihood." As above, we take as given that our observations of accident severity are accurate and are generated by the exponential distribution. Further, we accept the Reactor Safety Study estimate with one modification. The modification is that we will change the error parameter ("approximate uncertainty for nuclear events" in footnote to Figure 1 above) for the




FIGURE 1. ALTERNATIVE ESTIMATES OF FREQUENCY OF NUCLEAR ACCIDENTS FROM ACTUAL EXPERIENCE AND REACTOR SAFETY STUDY

estimates from the Reactor Safety Study estimate by a confidence factor,  $\lambda$ . As noted above, the Reactor Safety Study states that the risk estimates have a possible error factor of 5 either way. If we accept this error parameter, we have a confidence factor  $\lambda = 1$ . If we feel that the Reactor Safety Study underestimates the error parameter for its risk estimates by a factor of 2, then we have a confidence factor of  $\lambda = 1/2$ . If we feel the study is worthless, the confidence factor is  $\lambda = 0$ . A more detailed description of the technique is given in the appendix.

The most likely estimate of the risk of nuclear power accidents is given in Table 2. Recall that these estimates use the technique of "maximum likelihood"--asking what is the most likely frequency of severe nuclear accidents given the Reactor Safety Study and past history.

Not surprisingly, the results indicate that our revised estimates of the frequency of severe accidents have changed dramatically. If our faith in the Reactor Safety Study is unshaken ( $\lambda = 1$ ), and we simply think we have had very bad luck, our estimate of the frequency of severe accidents would be revised upward from about 32 per million reactor years (MRY) to 1250 per MYR. At the other extreme, if we think the Reactor Safety Study methodology is fundamentally flawed and therefore worthless, our estimate rises from 32 per MRY to 2500 per MRY. It is somewhat surprising to note that the degree of confidence in the Reactor Safety Study has little effect on the revised risk parameter.

Table 2. Alternative estimates of risk parameter with different levels of confidence in Reactor Safety Study

(1) Confidence parameter ( $\lambda$ )	(2) Revised risk estimate of accident frequency ( $\alpha$ )	(3) Upper and lower bound on risk estimate ( $\alpha^l, \alpha^u$ )
1	1250 per MRY	(6300, 160) per MRY
1/2	2000 per MRY	(7900, 200) per MRY
0	2500 per MRY	(10,000, 250) per MRY
Memorandum:		
Original Reactor Safety Study Estimate	32 per MRY	(160, 6) per MRY

Note: Confidence parameter ( $\lambda$ ) in column (1) is a subjective variable which is the ratio of error of the risk estimate in the Reactor Safety Study to the true error. The revised risk estimate in column (2) is the maximum likelihood estimate of the risk parameter ( $\alpha$ ) for severity of one death or more. The upper and lower bounds on the risk estimates ( $\alpha^l$  and  $\alpha^u$ ) are the 95% confidence interval value of the risk parameters from the maximum likelihood estimate for a lognormal distribution.

As a final note, we might ask how many years of accident-free experience in the future will be required before we again accept the Reactor Safety Study estimates of risk? With 100 operating reactors, the study estimates that an accident causing one or more fatalities should occur every 300 years. Thus we would have to wait until almost the year 2300 to vindicate the original estimates of the Reactor Safety Study. This indicates the extent to which the occurrence of "almost impossible" events can change our beliefs in almost irreversible ways.

#### 4. Summary

Our conclusions are simple:

1. With less than 100 nuclear power reactors operating in the United States, the Reactor Safety Study predicts an accident as severe as Three Mile Island should have occurred once every four hundred rather than once in less than ten years of intensive use of nuclear power. The chance of such a large discrepancy is so small that we must have very serious doubts about the reliability of the "macro" estimates of the risk of nuclear power in the Reactor Safety Study.

2. Combining the Reactor Safety Study and actual history, we obtain a revised estimate that (with 100 reactors) an event as severe as Three Mile Island will occur every 1.6 to 100 years, with the best guess being about every 20 years.

These results must be viewed in perspective and with caution for the following reasons:

3. The results are based on an unusual method of combining theory and history; other methods might yield different answers. In addition, many of the data on the effects of the Three Mile Island accident are preliminary and may well be revised by the work of the Kemeny Commission.

4. The results of this note do not attempt to ask why the Reactor Safety Study may have erred in its estimates. Indeed, from the "macro-analytical" perspective used here, as well as from a policy perspective, the source of error matters little. This point is of critical importance and will be elaborated. The introduction to the Report states:<sup>6</sup>

The Reactor Safety Study was sponsored by the U. S. Atomic Energy Commission to estimate the public risks that could be involved in potential accidents in commercial nuclear power plants of the type now in use.

<sup>6</sup>Reactor Safety Study, p. 1.

The errors from the study arise from two sources. First, the actual events analyzed may have been analyzed incorrectly. Second, certain classes of events may have been overlooked or omitted. Preliminary accounts suggest that the importance of "human error" may have been underestimated, although human error was studied in great depth.

In any case, it is no excuse for a study with the purpose cited above that a certain class of events was omitted (implicitly setting the probability equal to zero). Rather, this is a grave methodological flaw. It is akin to asking for the frequency of homicides omitting considerations of human passion.

Further study will probably determine the extent to which Three Mile Island was due to errors of type one or type two above. In the first case, the study was wrong; in the second, irrelevant. In either case, the risks of nuclear power will have to be revised substantially.

5. Whatever the outcome of a new assessment of the risks of nuclear power, it would be a serious mistake to assume that if the Reactor Safety Study was too optimistic this implies nuclear power is an unviable energy option. Even with the revised probabilities suggested above, nuclear power has a smaller health effect per unit delivered energy than using coal. The risks are, however, no longer negligible.

Appendix

This appendix explains the technical steps needed to answer the question, "What is the most likely shape of the risk-severity function given both history and theory and depending on how we weigh the two?" The likelihood of the observed data on reactor operations can be written as:

$$L(\alpha; x, \lambda) = R(\alpha, \lambda) \cdot H(x, \alpha)$$

In this,  $x$  is the actual history observed to date (one accident in 400 reactor-years of operation);  $\alpha$  is the estimate of risk;  $\lambda$  is the estimate of the confidence of the result in the Reactor Safety Study discussed above.  $R(\alpha, \lambda)$  is probability of the parameter  $\alpha$  given the distribution found by Rasmussen;  $H(x, \alpha)$  is the likelihood of such a parameter given the observations from history;  $L(\alpha; x, \lambda)$  is the overall likelihood as the product of the two probabilities.

The specific technique is as follows: We assume that actual accidents of severity  $D \geq 1$  occur according to the exponential distribution  $f(T) = \alpha e^{-\alpha T}$ , where  $T$  is the number of reactor years of operation. The Rasmussen estimate of  $\hat{\alpha} = 32 \times 10^{-6}$  per reactor year is assumed to be lognormally distributed with standard deviation  $s$  estimated to be  $1.6 = \ln 5$  if  $\lambda = 1$ . The standard deviation is then divided by  $\lambda$  as our confidence decreases.

Under these assumptions, the likelihood of the historical observation is

$$L(\alpha; x, \lambda) = \alpha \exp(-T\alpha) (2\pi\lambda^{-2} s^2)^{-1/2} \exp\left[-\frac{(\ln \alpha - \ln \hat{\alpha})^2 \lambda^2}{2s^2}\right]$$

The maximum likelihood estimates in Table 2 are calculated as the value of  $\alpha$  which maximizes  $L$  for alternative values of  $\lambda$ .

In addition, we calculate an upper ( $\alpha^u$ ) and lower ( $\alpha^l$ ) bound on the risk estimate as ones where the ratio of the likelihood to the maximum for given value of  $\lambda$  is equal to 4.4. For the lognormal distribution, the parameter  $\alpha$  will lie between the two limits ( $\alpha^u$  and  $\alpha^l$ ) with 95 percent confidence.