

Comments on and Analysis of 1957 Sandia Report  
“Acceptable Military Risks from Accidental Detonation of Atomic Weapons”  
Martin E. Hellman, Stanford University

**CAVEAT:** I do not necessarily agree with the basic assumption of this report, that an acceptable level of risk for an accidental detonation of an atomic weapon would be such that the expected number of American lives lost per year would be comparable the number lost in natural disasters such as hurricanes, tornadoes, and earthquakes. My use of the report’s terminology therefore should not be confused with my own beliefs.

When I read Eric Schlosser’s acclaimed [2013 book](#), *Command and Control: Nuclear Weapons, the Damascus Accident, and the Illusion of Safety*, I found a tantalizing revelation on pages 170-171 (emphasis added):

THE ENGINEERS AT SANDIA knew that nuclear weapons could never be made perfectly safe. Oskar Morgenstern— an eminent Princeton economist, military strategist, and Pentagon adviser— noted the futility of seeking that goal. “Some day there will be an accidental explosion of a nuclear weapon,” Morgenstern wrote. “The human mind cannot construct something that is infallible . . . the laws of probability virtually guarantee such an accident.” . . . In order to reduce the danger, weapon designers and military officials wrestled with two difficult but interconnected questions: ***What was the “acceptable” probability of an accidental nuclear explosion?*** And what were the technical means to keep the odds as low as possible?

The Army’s Office of Special Weapons Developments had addressed the first question in a 1955 report, “Acceptable Military Risks from Accidental Detonation of Atomic Weapons.” It looked at the frequency of natural disasters in the United States during the previous fifty years, quantified their harmful effects according to property damage and loss of life— and then argued that accidental nuclear explosions should be permitted on American soil at the same rate as similarly devastating earthquakes, floods, and tornadoes. According to that formula, the Army suggested that the acceptable probability of a hydrogen bomb detonating within the United States should be 1 in 100,000 during the course of a year. The acceptable risk of an atomic bomb going off was set at 1 in 125.

I have [been working](#) on the related questions: What is an acceptable probability for nuclear deterrence failing, and what is the current level of risk compared to that value? Since almost no prior work existed for these or related questions, Schlosser’s revelation that this 1955 report existed caught my attention. On page 527, in his notes, Schlosser states:

Although I did not obtain the [1955] Army study, its conclusions are explored in “Acceptable Premature Probabilities for Nuclear Weapons,” Headquarters Field

Command, Armed Forces Special Weapons Project, FC/ 10570136, October 1, 1957 (SECRET/ RESTRICTED DATA/ declassified).

Unable to find either the 1955 or 1957 reports on line, I communicated with Mr. Schlosser, who kindly sent me the [1957 report](#) and [DTRA's response](#) to his FOIA inquiry, along with permission to place them online – just click on those links. The only changes I made were to remove his address from the FOIA response and to run OCR on the report so it is somewhat searchable. (The OCR was far from perfect.)

Filing FOIA requests are time consuming, and Mr. Schlosser deserves the gratitude of the research community both for pursuing this request and for sharing the results with the larger community.

Here's a summary of what I gleaned from the 1957 report:

Table I on PDF page 5 gives the 1955 study's conclusion as to acceptable rates of accidental detonation versus yield. This 1957 report never says explicitly that the risk of an accidental nuclear explosion would be acceptable if the expected number of deaths per year from such an event were the same as the average number of American deaths from natural disasters. But the earlier, 1955 report on which it is based may. In my reading, here's all I found:

A survey of available literature uncovered only one study that might furnish a logical basis for determining premature probability requirements (Reference 1).<sup>1</sup> This study ... attempted to relate accidental nuclear explosions to the disasters (rate and magnitude) which the nation has experienced in the past half century.  
...

[But,] whether or not the accident would be considered in the same light as a naturally occurring catastrophe (earthquake, flood, tornado, etc.) is problematical. In acts of nature, the tendency seems to be to treat these as tragic but unavoidable. In a nuclear disaster, there will likely be a tendency to blame the "irresponsible" military and scientists. As a result, public reaction could cause political pressure to be placed upon the military, thereby hampering or curtailing weapon programs. ...

In the aforementioned OSWD study, major U.S. catastrophes of the past 50 years were analyzed and an equivalent yield associated with each, based upon the people killed and/or property damage resulting from them. From this data, a frequency versus yield curve was drawn and the inference reached that since the nation has survived these, it could survive a similar rate in the future. ... Table I shows the proposed rates. Values are the acceptable number of weapon accident per year within each yield range shown. ...

---

<sup>1</sup> Reference 1 is the 1955 report.

The figures presented are debatable.

(a) During a full-scale war, the number of accidental detonations will be overshadowed by disasters resulting from enemy action. Therefore, it is doubtful whether there should be a prescribed rate [for wartime] based on peacetime safety considerations. The acceptable wartime rate should be that which would not significantly hamper our operational capability.

(b) The conversion from “normal” peacetime disasters to equivalent kiloton nuclear accidents neglects public reaction. ...

Table I ... assumed that an accident would result in a full-scale nuclear yield of one weapon only.

With this information, I went back to Schlosser’s statement from the above excerpt that:

... the Army suggested that the acceptable probability of a hydrogen bomb detonating within the United States should be 1 in 100,000 during the course of a year. The acceptable risk of an atomic bomb going off was set at 1 in 125.

I now could see that he associated “hydrogen bomb” with the report’s limits on a warhead having a yield greater than 10 MT and “atomic bomb” with one having a yield in the 1-10 KT range. As can be seen from Table I (on page 5 of the PDF), the report had 1E-5 (1 in 100,000) as the acceptable number of peacetime accidents per year involving warheads >10 MT, and 8E-3 (0.008 = 1 in 125) for those with yields in the 1-10 KT range. Those are the numbers in the above excerpt from Schlosser’s book.

The report says almost nothing about how it came up with Table I, so I tried to figure out the assumptions it used. For the moment assuming that the first entry in the Table is correct (an acceptable annual risk of 1E-5 for a premature detonation with a yield of 10 MT), how did they come up with the later entries (peacetime only)?

I guessed that Table I used a simple model in which each successive row in the table should have an acceptable annual probability which is  $10^{(2/3)} = 4.64$  times as large as the one above it.<sup>2</sup> The table below shows good agreement between my assumed model (the third column) and Table I’s entries (the second column).

While there is good agreement, there is enough error to indicate slightly sloppy calculations on the part of the report.

---

<sup>2</sup> In a simple model, the volume affected by an explosion grows as the cube of the radius, so the radius for a given level of destruction grows as the yield to the one-third power. Since the area affected grows as the radius squared, it (and presumably deaths and property damage) would grow as the two-thirds power of the yield. Equivalently, the acceptable annual risk would have to decrease as the two-thirds power of the yield.

To complete the picture of how Table I was constructed, we need to know how it came up with the first row's peacetime entry of 1E-5 (1 in 100,000). The report does not say how it derived this number, but it seems pretty clear that it intended to allow the same number of expected American deaths per year from an accidental nuclear detonation as from natural catastrophes, such as earthquakes, floods, hurricanes, and tornadoes.

Yield	Acceptable Annual Risk per Table I	Acceptable Annual Risk Relative to First Entry
>10 MT	1.00E-05	1.00E-05
1-10 MT	5.00E-05	4.64E-05
100-1000 KT	3.00E-04	2.15E-04
10-100 KT	2.00E-03	1.00E-03
1-10 KT	8.00E-03	4.64E-03
0.1-1 KT	5.00E-02	2.15E-02

On page 4 of the PDF, the report notes that the acceptable probabilities of premature detonation were based on data for the previous 50 years:

A survey of available literature uncovered only one study that might furnish a logical basis for determining premature probability requirements (Reference 1). This study ... attempted to relate accidental nuclear explosions to the disasters (rate and magnitude) which the nation has experienced in the past half century.

I don't have the same data set, but found that, more recently, American deaths due to weather-related natural disasters (flood, lightning, tornado, hurricane, etc.) [averaged 527 per year](#), and weather would seem to be responsible for most American deaths due to natural disasters. During the 50 year period presumably used by the report (1907-1957), deaths to earthquakes (the only non-weather related natural catastrophe mentioned in the report, or that I could think of as being significant) [totaled 342](#), which is less than 7 per year. If we go back one year earlier, the 1906 San Francisco earthquake and fire would have been included and added 3,000 deaths. But even that would only increase the 50-year average number of deaths per year due to earthquakes to 67 per year – still small compared to roughly 500 per year due to weather.

In conclusion, I used 500 American deaths per year as being the average due to natural disasters, rounding 536 down to avoid giving the impression of greater accuracy than is warranted with such approximate estimates. (If the San Francisco earthquake is included, I should use 600 deaths per year average instead.)

The first row of Table I, with an acceptable probability of 1E-5, then would seem to indicate that a >10 MT premature nuclear explosion was expected to kill  $500 \times 1E5 = 50$

million Americans – roughly a third of the US population at that time, and seemingly a ridiculously high figure. (If the San Francisco earthquake is included, that would result in 60 million, an even more ridiculous figure.)

One possible partial explanation for this apparent error is that the number of deaths due to natural disasters may have been lower in their 50-year period (presumably from around 1900 to 1950) than in recent years. But, even halving the number, to 250, would still seem to imply that they thought 25 million Americans (roughly one-sixth of the population) would be killed by a premature explosion. I also suspect that, while the US population was smaller during their 50 year period, poorer flood control, emergency response, etc. caused the actual number of deaths per year to be higher than in recent years.

Maybe the report was using a worst case situation, as opposed to an average one (as in the rest of their analysis)? Or is this indicative of sloppy work? My earlier analysis of Table I also indicated possibly sloppy work.

Next I turned to Table III on page 9 of the PDF. At first it confused me because it claims to answer the question, “What assurance is given that there will be no nuclear accidents?” It doesn’t say over what period of time, and relies on Table II (redacted, presumably because it would divulge the number of weapons in each yield range), whereas it seems to me that Table I gives the answer directly for a one year period.

My friend and colleague, Dr. Richard Duda, figured out the explanation, which was not immediately evident from the report: it assumed a 10 year period. Once that missing data point is added, and the Poisson assumption (mentioned as a footnote to Table III) is used, we get the following table. The third column is the “assurance” to three decimal places, while the last column is the entry given in Table III. It is seen that they did not round, but used 0.99 for any number greater than or equal to 0.99 and less than 1.

Yield	Annual Probability of Premature Detonation per Table I	Probability of no Premature Detonations in 10 years (Poisson assumption)	Table III Entry
>10 MT	1.00E-05	1.000	0.99
1-10 MT	5.00E-05	1.000	0.99
100-1000 KT	3.00E-04	0.997	0.99
10-100 KT	2.00E-03	0.980	0.98
1-10 KT	8.00E-03	0.923	0.92
0.1-1 KT	5.00E-02	0.607	0.60

While this explains how Table III was computed, it leaves open the question of why 10 years was deemed an acceptable period over which to give assurance at different levels that no premature nuclear detonation would take place. Both Dr. Duda and I thought 10 years was far too short a period of time. If we are right, that's one more indication of sloppy work.

Temporarily forgetting about the sloppy work, and whether the report is correct in allowing the expected number of American deaths per year due to a nuclear failure to be comparable to the number killed in natural disasters, here is what the same basic methodology would predict is an acceptable level of risk for nuclear deterrence failing in two different modes:

A failure of nuclear deterrence which results in a nuclear terrorist incident destroying part of an American city such as New York is likely to kill on the order of 100,000 Americans. Using 500 American deaths per year from natural disasters results in a required annual risk of such a nuclear terrorist event of at most 1 in 200 or 0.5%. Equivalently, we would have to expect such an event not to happen for approximately 200 years. Contrasting this with Henry Kissinger's estimate of 10 years in the *Nuclear Tipping Point* documentary, or Dr. Richard Garwin's [congressional testimony](#) estimating the risk at 10-20% per year, we see that if those experts are right, we need to reduce the risk by more than a factor of 10. While both of those estimates are subjective, they come from men with significant expertise in matters of national security.

Similarly, if nuclear deterrence suffers a complete failure, resulting in an all-out nuclear exchange, and if that is assumed to kill even 100 million Americans (1,000 times as many as in the assumed nuclear terrorist incident), then the risk of that event would have to be 1-in-200,000 per year (1,000 times smaller than for the nuclear terrorist incident). Equivalently, we would have to expect such an event not to happen for approximately 200,000 years. Again, this seems at least an order of magnitude lower than any reasonable estimate of the current level of risk.

Some have argued that a full-scale nuclear war might kill all of humanity, in which case an even lower probability of failure would be required, even using the methodology of the 1957 report. Some have also argued, either on moral grounds or on the cost to unborn generations, that the cost of human extinction cannot be measured on the basis of the expected number of lives lost. Such arguments would predict an even lower acceptable probability of failure. It is instructive to note, however, that even using what some would see as the overly optimistic methodology of the report, the acceptable failure rate is much lower than most people's subjective estimate: When I ask people for an order of magnitude estimate for how long they think nuclear deterrence will work before failing and causing a full-scale nuclear war, the vast majority see 10 years as too short and 1,000 years as too long, leaving 100 years as their order of magnitude, subjective estimate.