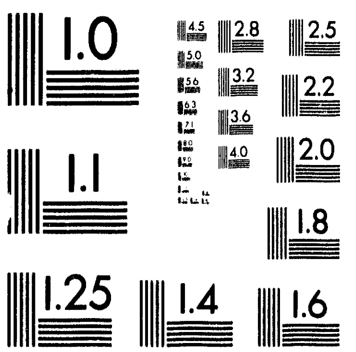
1.0
1.1
1.25
1.4
1.6
1.8
2.0
2.2
2.5
2.8
3.2
3.6
4.0
4.5
5.0
5.6
6.3

Conf-940312--12

LA-UR 93-3461

Los Alamos National Laboratory is operated by the University of California for the United States Department of Energy under contract W-7405-ENG-36

TITLE: IS PROBABILITY OF FREQUENCY TOO NARROW?

AUTHOR(S): Harry F. Martz

SUBMITTED TO: To be presented at the PSAM-II Conference, March 20-25, 1994, San Diego, CA

Los Alamos

Los Alamos National Laboratory
Los Alamos, New Mexico 87545

IS PROBABILITY OF FREQUENCY TOO NARROW?

Harry F. Martz

Group A-1, MS F600
Los Alamos National Laboratory
Los Alamos, NM 87545

Modern methods of statistical data analysis, such as empirical and hierarchical Bayesian methods, should find increasing use in future Probabilistic Risk Assessment (PRA) applications. In addition, there will be a more formalized use of expert judgment in future PRAs. These methods require an extension of the probabilistic framework of PRA, in particular, the popular notion of probability of frequency, to consideration of frequency of frequency, frequency of probability, and probability of probability. The genesis, interpretation, and examples of these three extended notions are discussed.

FREQUENTIST AND SUBJECTIVIST PROBABILITY

In PRA theory and practice, there are two basic interpretations of probability: *frequentist* (or objectivist) and *subjectivist* (or Bayesian). The distinguishing notion of the frequentist interpretation is that probability is defined as the *limiting relative frequency* of occurrence of a random event in an infinite sequence of identical trials. On the other hand, the subjectivist view considers probability to be a subjective statement of an analyst's personal *state-of-knowledge or degree-of-belief* regarding the occurrence of an event based on the available evidence. It is generally believed that both interpretations are appropriate for use in PRA. Parry (1986) and others argue that it is important to separate the effects of both classes of uncertainty on results typically of interest in PRA.

Apostolakis (1990) refers to the model under consideration as the *model of the world.* For example, in Level 1 PRA, the model of the world consists of the event trees, fault trees, etc. used to represent the response of a facility to a particular accident.

For purposes here, consider two basic classes of uncertainties in the model of the world. The first class are those uncertainties which are virtually impossible (either technically or economically or both) to either reduce or remove. Collectively denote the set of such irreducible uncertainties as *stochastic or random uncertainties*. For example, variability in the quality of manufactured materials, uncertainties in equipment performance, and operator errors are typical examples of random uncertainties. Various well-known probability distributions (the formal name for the simplest type of probability models of the world), such as the normal, binomial, or Poisson distributions, are commonly used to model basic random uncertainties. These probability distributions involve various parameters (such as the failure rate or intensity in the Poisson distribution) which are usually unknown and must be estimated. The frequentist interpretation of probability is generally used to express random uncertainties.

The second basic class of uncertainties are *state-of-knowledge* (also known as *systematic*) *uncertainties*. These uncertainties represent our lack of precise knowledge about uncertain parameters in our model of the world (or even uncertainty regarding our overall choice of model of the world). For example, we may not know the precise value of the failure rate parameter in an exponential distribution used to model the time-to-failure of an operating pump. Based on the available data and experience with such pumps, we represent our uncertainty about this failure rate in the form of a probability distribution. Because this distribution represents our current state-of-knowledge about the unknown failure rate parameter (without any regard for consideration of repeated trials), it embodies a subjectivist interpretation of probability.

In contrast to irreducible random uncertainties, it is possible to reduce state-of-knowledge uncertainties by improving or enhancing our state-of-knowledge. In the limiting case, we have perfect state-of-knowledge and the corresponding probability distribution would be a delta function. We can improve (update) our state-of-knowledge by obtaining additional pertinent data or information. Bayes' theorem is appropriate for updating our state-of-knowledge distribution given additional data or information, a process known as *Bayesian updating*. Because there are often alternative models for certain phenomena of interest in PRA, the uncertainty inherent in subjectively choosing a model (so-called *modeling uncertainty*) is a state-of-knowledge uncertainty.

Subjective probability is also an appropriate way to express uncertainty about events, the outcome of which cannot easily be interpreted as an imagined series of repetitions. Such events are often referred to as *propositions*. For example, we may be interested in the event "a manned mission to Mars occurs prior to the year 2050." This is certainly an uncertain event (as of 1993), and a repetitive series of trials of this event is difficult to conceive because it involves a specific (unique) period of time which will occur only once. Thus, it is appropriate to interpret our uncertainty about such events using subjective probability. Notice, however, that determining whether or not a particular event can be imagined as a repetitive series of trials is also often quite subjective, especially if we permit the use of "thought" experiments. While one person may be able (and willing) to entertain such an imagined series of trials, another person may find an imagined series of identical repetitions to be totally inconceivable. Thus, the most appropriate notion of probability to use to describe an uncertain event is also sometimes a subjective decision.

The bottom-line quantities of interest in a PRA are usually either uncertain exceedence probabilities or frequencies of occurrence per year (such as the probability with which damage level x or greater occurs per year or the frequency of occurrence of severe core damage per year). Clarotti (1988) and Howard (1988) argue that the frequency of occurrence of rare events considered in PRA are fictional parameters which should not be interpreted within the context of a repetitive series of trials. Thus, there is some disagreement regarding whether the frequency of occurrence of rare events per unit time should be interpreted using relative frequency or degree-of-belief.

To distinguish between the two notions of probability, I shall denote frequentist (relative frequency-based) probability as probability_f (where the subscript "f" stands for "frequentist") and subjectivist (degree-of-belief-based) probability as probability_s (where the subscript "s" denotes "subjectivist").

PROBABILITY OF FREQUENCY

Kaplan and Garrick (1981) introduced the notion and use of *probability of frequency* in PRA. This approach is further described in Kaplan et al (1981). They consider the relative frequency of occurrence of an event of interest (that is, the fraction of times that it occurs) in a series of identical repeated trials in which the event has a constant probability_f of occurrence p on each trial. They denote this relative frequency as ϕ. Of course, by definition, ϕ approaches p in the limit as the number of repeated future trials approaches infinity. Thus, interest in ϕ (over an infinite series of trials) is equivalent to interest in p.

Because of inherent parameter and modeling uncertainties in our model of the world, we are usually uncertain about the precise value of p (especially so in the case of rare events). By performing an *uncertainty analysis*, we can express our state-of-knowledge about the unknown value of p in the form of a probability_s distribution (or curve) over p_f. Kaplan and Garrick refer to these as *probability_s of probability_f* (or *probability_s of frequency*) curves. The outer probability_s is interpreted as *Bayesian confidence* (*confidence_s* or *credibility*). For example, using the probability_s of frequency approach, we can make statements such as "we are 95% confident_s that the annual probability_f of 10 or more early fatalities resulting from internally initiated accidents at Plant X is not larger than 10^{-7}".

FREQUENCY OF FREQUENCY

Empirical Bayes (EB) methods are a class of modern statistical inference techniques which are potentially useful in PRA applications for estimating unknown parameters in the model of the world. To illustrate the use of EB in PRA, Martz and Abramson (1994) use EB methods to estimate the average (over time-and-station) probability_f that emergency diesel generators (EDGs) at US commercial nuclear power plants successfully load-run on a given demand (the so-called *load-run demand reliability*). For simplicity, denote the time-and-station average of this load-run demand reliability as p. Martz and Abramson are interested in estimating p using data from a large number of stations. EB methods are appropriate for use in those situations where a parameter, such as p, is not constant over a corresponding population of interest (in this case, the population of stations). EB methods assume that p varies randomly over the population according to some unknown *EB prior distribution g(p)* [Maritz and Lwin (1989); Martz and Waller (1991)]. The EB prior g(p) expresses the probability_f distribution of p over the EDG population and, for this reason, is sometimes referred to as a *population variability distribution* (or curve). In our example, g(p) expresses the fraction (or relative frequency) of stations whose average EDG load-run demand reliability falls in each reliability increment.

It is important to note that, in spite of the reference to Bayes in the name "EB", EB methods make no use of degree-of-belief (that is, probability_S). The reference to Bayes is due to two aspects of EB: (1) the notion and use of a "prior distribution" for p, and (2) the use of Bayes' theorem as an integral part of the EB procedure. The advantage of EB over Bayesian methods (when both can be used) is that EB methods involve less subjectivity, and thus are more "data driven", than pure Bayesian methods.

In a plant-specific PRA, we are not primarily interested in the overall population of stations; rather, we are interested in the EDG reliability of a particular station. Bayes' theorem is the tool for "specializing" (by updating) the EB prior on p to the particular value of p, say p_i, for a particular station (the i^{th} here). As in Bayesian estimation, we determine an EB posterior distribution for p_i. Unlike a Bayesian posterior distribution, this station-specific EB posterior distribution embodies probability_f (as does the EB prior), and is a *frequency of frequency* distribution.

What interpretation do we give to this station-specific EB posterior distribution? Suppose that, for our particular station, we have observed $x_i = 2$ load-run failures in $n_i = 117$ total station EDG demands. The EB posterior distribution of p given $x_i = 2$ and $n_i = 149$ [that is, $g(p|x_i=2; n_i=117)$] expresses the fraction (or relative frequency) of stations in the population for which 2 failures are observed in 117 demands whose load-run demand reliability falls in each reliability increment. Because we have actually observed 2 failures in 117 demands for our particular station, this is the appropriate station-specific probability_f distribution that we use to make inferences (such as point and interval estimates) about p_i.

FREQUENCY OF PROBABILITY

Consider a one-time event which either will or will not occur -- except that we don't know which -- and for which a repetitive series of trials is inappropriate or impossible to consider. For example, consider again the event that "a manned mission to Mars occurs prior to the year 2050." Now, this future event will certainly either occur or it won't occur-- we just don't know which. In PRA parlance, this is known as a *zero-one* event.

Kaplan et al (1981) liken this situation to that in which either a two-headed or a two-tailed coin is to be flipped, and we just don't know which coin it is. Anyone can then express his or her state-of-knowledge (degree-of-belief) regarding the occurrence of the Mars event as the probability_S p that the two-faced coin is two-headed. Further, suppose that we are interested in assessing the variability of p_i over some population of interest; for example, the population of Los Alamos National Laboratory employees. If we now ask each Los Alamos employee for his or her value of p (and everyone complies), then we can construct a population variability curve that expresses the fraction (or relative frequency) of employees at Los Alamos whose value of p falls in each increment. Again, using the Kaplan and Garrick nomenclature, this population variability curve on p is a *frequency of* $\textit{probability}_S$ distribution.

Figure 1 shows the frequency distribution for the probability_S p of the Mars event (in the form of a relative frequency histogram) based on a sample of size n = 56 Los Alamos employees. It is interesting that this distribution tends to be trimodal, having a delta function at 0.5 (which appears in Figure 1 as a rectangle) and a "bump" on either side of (but distinctly away from) 0.5. The two bumps correspond to strongly optimistic or pessimistic employees, while those who responded with p = 0.5 are indicating that they no nothing (and are thus "guessing"), regarding the occurrence of the Mars event. For these data, we find that the mean is 0.49 and that $\text{Probability}_f\,(0.01 < p < 0.92) = 0.80$.

The formal use of expert opinion is an important part of modern PRA theory and application [Wheeler et al (1989)]. Suppose that we expand our PRA model of the world to include the representation of the variation in opinions of different experts. Within this expanded model, the notion of frequency of probability_S is appropriate. Let me explain. In PRA analysis, many issues can often best be interpreted as zero-one events in which an expert is asked to quantify his or her degree-of-belief regarding the occurrence of the event. The uncertainty is usually elicited in the form of a probability_S p. For example, in the reactor context, consider the event "in a proposed advanced reactor design, natural processes (such as passive heat removal) fail to successfully cool the reactor core." Suppose that we consider a panel of experts for this event to be a sample from a larger population of experts. In this case, frequency of probability_S is an appropriate way to express the variation in degree-of-belief within the population of experts.

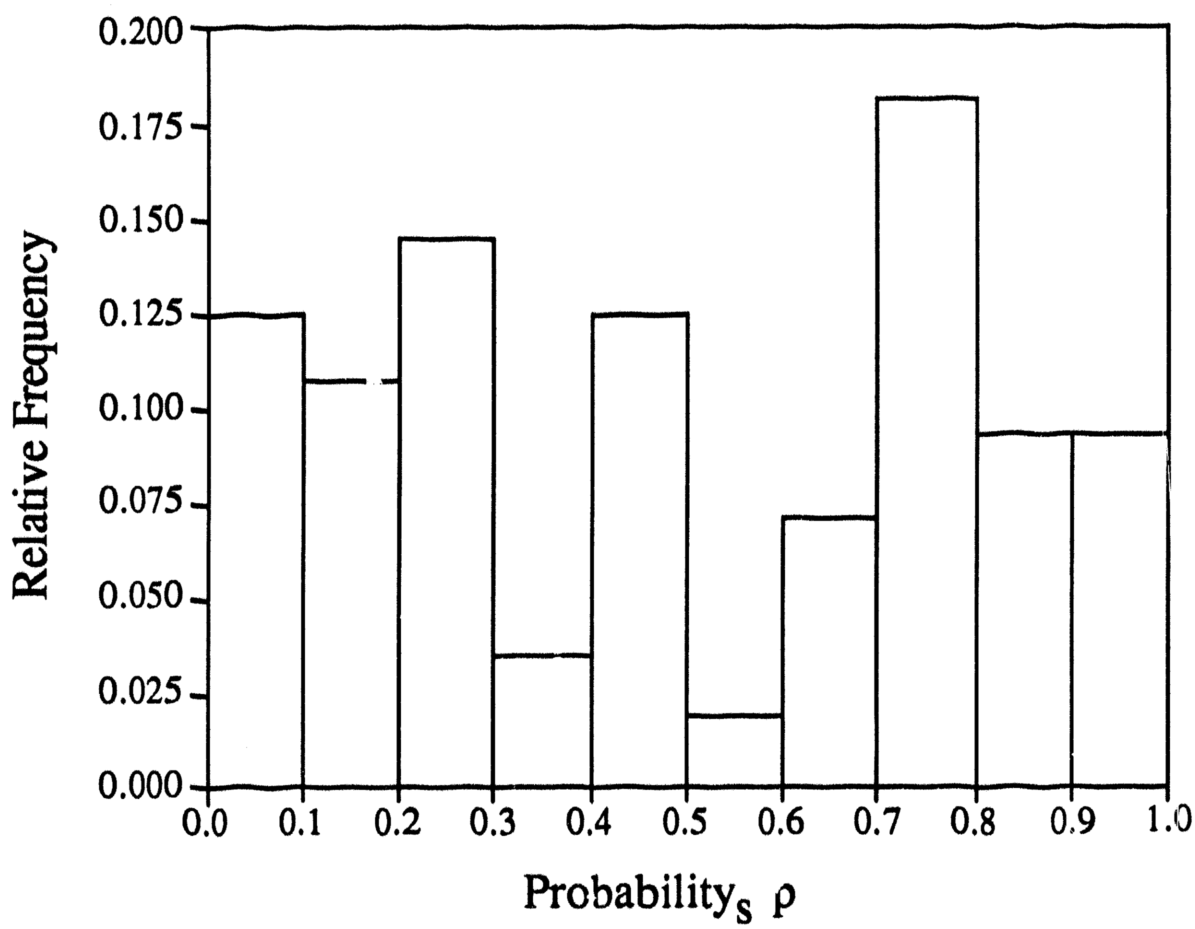

Figure 1. Frequency of probability distribution for the probability$_S$ p of a manned mission to Mars prior to the year 2050.

How would we update a frequency of probability$_S$ distribution as new information is obtained? We must be careful here to define what we mean by "new" information. Because the event of interest is zero-one, a single trial (if such would even be possible) is sufficient to completely remove all uncertainty about the event. This is not what I mean by new information. Rather, by "new information" I mean any change that an expert makes in his or her own subjective probability$_S$ p as a result of additional information or new knowledge. For example, in the Mars event, new knowledge might consist of an emerging global political and economic climate advocating that such a mission be undertaken. Each expert would then be asked to update his or her own value of p in light of this new information. Given the new value(s) of p, these estimates would similarly be used to determine a new (updated) frequency of $probability_S$ distribution in light of the new information. Note, however, that such updating is not performed using Bayes' theorem.

PROBABILITY OF PROBABILITY

Here we are basically concerned about what is believed about what is believed, and, at first glance, such a statement appears to be nonsense. If, however, I am basically concerned about expressing what an analyst believes about what an expert (or experts) believes (regarding the zero-one event), then this statement makes perfectly good sense. Further discussion regarding this statement is given by Pearl (1988), pp. 358-359.

As in the previous case, we now expand our PRA model of the world to include an expression of the variety and quality in the opinions of different experts as viewed by an analyst acting as a "super expert". Thus, within this expanded model, the notion of *$probability_S$ of $probability_S$* is defined. There are at least two possible interpretations of $probability_S$ of $probability_S$.

First, the analyst can represent his/her belief regarding the expert opinions in the form of a $probability_S$ distribution over the elicited probabilities. This situation is similar to the case of frequency of $probability_S$, except that now the outer distribution is a $probability_S$ distribution that must be subjectively assessed by the analyst. For example, the commonly used decision to equally weight each expert judgment is one such (outer) $probability_S$ distribution.

There is another interpretation which is also appropriate. Suppose that the analyst is interested in the elicited value from a particular expert, say Jane. Prior to eliciting Jane's value of p, p_{Jane}, the analyst may wish to express his/her personal belief about p_{Jane} in the form of a "prior" $probability_S$ distribution on p_{Jane}. This distribution expresses any substantive information that the analyst may have regarding Jane's ability in such tasks (including any known or perceived biases).

How do we update the analyst' prior once Jane has given her value? One possibility is to consider an appropriate subjective likelihood for Jane. This likelihood is a statistical model that relates the likelihood that Jane will report a value p_{Jane} relative to a true value of p. What do I mean by a "true value of p"? Subjective opinions being what they are, there clearly is no true value of p. However, one can define the true value of p to be the average value of p that would be elicited from a large population of experts. Examples of such conceptual likelihoods involving expert opinion are considered by Mosleh and Apostolakis (1984) and Martz and Bryson (1984). An important feature that is captured in such models is the consideration and inclusion of possible "biases" that are known to accompany expert opinions. Given the analyst' prior distribution and the assumed conceptual subjective likelihood model for Jane, Bayes' theorem is then used to obtain the desired updated posterior distribution for p_{Jane} given Jane's estimate.

CONCLUSIONS

Future PRA applications will likely use modern statistical methods, such as empirical and hierarchical Bayes methods, for estimating uncertain parameters in the PRA model of the world. In addition, the model of the world will be enlarged to include additional concerns involving the elicitation and combination of expert opinions. Because neither of these extensions are adequately considered using probability_S of frequency, I believe that this notion must be expanded to permit consideration of frequency of frequency, probability_S of frequency, and probability_S of probability_S.

Many authors have pointed out the importance of carefully distinguishing between random and state-of-knowledge uncertainties in PRA [Cornell (1984), Parry (1986), Apostolakis (1990)]. The primary concern in using these extended notions is how to separate the contributions from random and state-of-knowledge uncertainties in the results. It is not altogether clear that a clean separation is possible. If this task proves impossible (and well it might), perhaps the only reasonable alternative is to argue from a pure Bayesian perspective that all uncertainty, regardless of its origin, is really of the degree-of-belief type [Lindley (1985)].

Clearly there are other issues involved in using these extended notions which have not been considered here. For example, issues involving Monte Carlo simulation using these new notions are not discussed, and Bayesian updating has not been adequately addressed. However, I hope that this paper will stimulate additional interest in these and related issues.

ACKNOWLEDGMENTS

This work was supported by the US Nuclear Regulatory Commission, Office of Nuclear Regulatory Research, under contract FIN L1836-3 at Los Alamos National Laboratory. I gratefully acknowledge this support.

REFERENCES

Apostolakis, G., 1990, The concept of probability in safety assessments of technological systems, *Science* 250:1359.

Clarotti, C. A., 1988, *in* : "Reliability Engineering," A. Amendola and A. Saiz de Bustaman, eds., Kluwer Academic, Dordrecht, Netherlands, pp. 49-66.

Cornell, C. A., 1984, Seismic hazard analysis: some current directions ans some comments on its role in seismic PRA, presented at the *1984 Pressure Vessels and Piping Conference*, San Antonio, Texas, PVP-Vol. 79, p. 73, ASME.

Howard, R. A., 1988, Uncertainty about probability: a decision analysis perspective, *Risk Analysis* 8:91.

Kaplan, S. and Garrick, B. J., 1981, On the quantitative definition of risk, *Risk Analysis* 1:11.

Kaplan, S. et al, 1981, Methodology for probabilistic risk assessment of nuclear power plants, *Pickard, Lowe and Garrick, Inc. Report* PLG-0209, Irvine, California.

Lindley, D. V., 1985, "Making Decisions (2*nd* ed.)," John Wiley, New York.

Maritz, J. and Lwin, T., 1989, "Empirical Bayes Methods (2*nd* ed.)," Chapman and Hall, London.

Martz, H. F. and Abramson, L. R., 1994, Empirical Bayes estimation of the average reliability of nuclear power plant emergency diesel generators to load-run upon demand, presented at the *PSAM-II Conference*, San Diego, California.

Martz, H. F. and Bryson, M. C., 1984, A statistical model for combining expert opinions, *IEEE Trans. on Rel.* R-33:227.

Martz, H. F. and Waller, R. A., 1991, "Bayesian Reliability Analysis," Krieger, Malabar, Florida.

Mosleh, A. and Apostolakis, G., 1984, Models for the use of expert opinions, in: "Low-Probability/High-Consequence Risk Analysis," R. A. Waller and V. T. Covello, eds., Plenum, New York, pp. 107-124.

Parry, G. W., 1986, A discussion on the use of judgment in representing uncertainty in PRAs, *Nucl. Eng. Des.* 93:135.

Pearl, J., 1988, "Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference," Morgan Kaufmann, San Mateo, California.

DATE
FILMED
2 / 4 / 94
END