ON THE USE OF A BAYESIAN REASONING IN SAFETY AND RELIABILITY DECISIONS—THREE EXAMPLES

REACTOR SITING

KEYWORDS: *safety, reliability, decision-making, thermal power plants, transport, nuclear fuels, probability, steam turbines, failures, turbine blades, safety engineering*

STAN KAPLAN and B. JOHN GARRICK
Pickard, Lowe and Garrick, Inc., 2070 Business Center Drive Suite 125, Irvine, California 92715

Received October 24, 1978
Accepted for Publication February 13, 1979

Bayes' theorem is used to quantify the impact of "new evidence" in three energy-related decision problems. The first problem concerns the risk of radioactivity release during the railroad transport of spent nuclear fuel. This history of shipments thus far is shown to make it highly unlikely that the frequency of release is on the order of 10^{-3} or greater per shipment.

The second and third applications involve predicting the availability performance of new generations of turbine blades. Bayes' theorem is demonstrated as a means for incorporating in the prediction the limited operational data on the new blades along with the experience of the earlier generation and the knowledge of the design changes.

I. INTRODUCTION

Decisions made in the areas of safety and reliability often have great impact on the public and private welfare. They are also occasionally the subject of much controversy, often in large measure because the information on which these decisions must be based is incomplete, outdated, only partially relevant, and subject to varied interpretation.

In these situations, it is useful to have a systematic, rational procedure that can incorporate all the bits and pieces of evidence, all the partially relevant data, and anything else that is known, into a quantitative prediction of the safety or reliability of the system or plant in question. If this procedure also has a kind of convincingness about it that can move the various parties in the direction of consensus, so much the better.

The well-known theorem of Bayes can be the basis for procedures of this sort. Indeed, it is exactly for this kind of situation that Bayes' theorem was designed.[a]

The purpose of this paper is to give three real-life examples of the application of Bayes' theorem, two concerned with predicting the reliability of turbines for large power plants and one with predicting the safety of railroad transportation of spent nuclear fuel. We precede the presentation of these examples with a review of the meaning of the probability and a brief derivation of Bayes' theorem.

II. PROBABILITY

There is, and has been for a long time, much controversy among several schools of thought on probability, notably the so-called "frequentists" and the "subjectivists." In this paper, we would like to adopt the attitude that there is no "right" definition of probability. We need simply to recognize that there are two different notions here (at least), and if we give each the dignity of its own name, we can avoid further miscommunication. Thus, for the purposes of this paper, when we are talking about the measurable results of an experiment, flipping coins or counting seedlings, etc., we use the work "frequency." The subject that deals with such measurements and the interpretation of such data we call "statistics." When we are talking about a state of knowledge, a state of confidence (which might derive from such measurements), we use the term "probability." The science of such states of confidence and

[a]While the theorem itself is well known, there is considerable misunderstanding of and controversy about its applicability to situations of this sort (see, for example, Refs. 1 and 2). We hope that the examples given here clear up some of the controversy.

how they rationally change with new information is what in this paper we mean by the "theory of probability."

The best expression of this definition of probability that we know of was given by E. T. Jaynes in a short course at the University of California, Los Angeles in 1960. We reproduce it from Jaynes' notes as follows:

> "Probability theory is an extension of logic, which describes the inductive reasoning of an idealized being who represents degrees of plausibility by real numbers. The numerical value of any probability (A/B) will in general depend not only on A and B, but also on the entire background of other propositions that this being is taking into account. A probability assignment is 'subjective' in the sense that it describes a state of knowledge rather than any property of the 'real' world; but it is completely 'objective' in the sense that it is independent of the personality of the user; two beings faced with the same total background of knowledge must assign the same probabilities."

To summarize then, probability is a numerical expression of a state of confidence, a state of knowledge—which may be influenced by statistical data, frequency measurements, if available. If not, then it is the lack of data that influences our state of knowledge and that is a reflected in the numerical values we assign.

III. BAYES' THEOREM

In basic probability theory, $P(A)$ is used to represent the probability of the occurrence of event A and, similarly, $P(B)$ would represent the probability of event B. To represent the joint probability of A and B, we use $P(A \wedge B)$, which is the probability of the occurrence of both event A and event B. Finally, the conditional probability is $P(A/B)$, which represents the probability of event A given that event B has occurred.

From a basic axiom of probability theory, the probability of the two simultaneous propositions, A and B, is given by

$$p(A \wedge B) = p(A)\, p(B/A) \ , \qquad (1)$$

and equivalently by

$$p(A \wedge B) = p(B)\, p(A/B) \ . \qquad (2)$$

Therefore, equating the right sides of Eqs. (1) and (2) and dividing by $p(B)$, we have

Bayes' Theorem:

$$p(A/B) = p(A) \left[\frac{p(B/A)}{p(B)}\right] \ , \qquad (3)$$

which says that $p(A/B)$, the probability of A given information B, is equal to $p(A)$, the probability of A prior to having information B, times the correction factor given in brackets.

This theorem, as powerful as it is simple, thus shows us how our probability, i.e., our state of confidence with respect to A, rationally changes upon getting a new piece of information.

Notice that in the language we have chosen to describe this theorem, we have already adopted the point of view set out in Sec. II regarding the meaning of probability. This point of view is the appropriate and useful one for the three examples that follow.

IV. EXAMPLE 1—RAILROAD TRANSPORT OF SPENT NUCLEAR FUEL

This example derives from several federal and state hearings involving the desire of several railroad companies to require that spent nuclear fuel be transported in so-called "special trains" that consist of an engine, a caboose, and the spent fuel car, and that are therefore much more expensive than an ordinary shipment.

In support of this desire, the railroads sought to show that a shipment of spent fuel is highly dangerous and in support of that, argued that the record of safe shipments so far achieved is of no significance and should have no effect on our state of confidence.

In the course of this line of argument, it was acknowledged that there have been, so far, ~4000 shipments of spent fuel in this country without a single release of radioactive material. However, the railroad's position was to dismiss this experience, pointing out that releases are expected to occur once in 10^8 or 10^{10} shipments, and that since 4000 is very small compared to 10^8 or 10^{10}, we are in a position of "poor actuarial statistics" and will not have enough statistics to improve that position in any reasonably foreseeable future.

From a Bayesian viewpoint, the 4000 release-free shipments should not be dismissed. They constitute a very important piece of evidence. While it is true that 4000 is infinitesimal compared to 10^8 or 10^{10}, and certainly true that in our lifetimes we will not amass anywhere near enough statistics to distinguish between release frequencies of 10^{-8} and 10^{-10}, that does not mean we are doomed to a position of poor statistics, and that the shipping experience we do amass will have no evidentiary value. Exactly the contrary is true. The point is that we do not really care whether the release frequency is going to be 10^{-8} or 10^{-10}. We want to know whether we can ship our spent fuel safely. That is, we want to know whether the release frequency is going to be on the order of 10^{-8} to 10^{-10} or on the order 10^{-3} to 10^{-4}. With respect to that question, our 4000 release-free shipments are very important evidence indeed, as we now show.

Let us suppose that we did not have the evidence of the 4000 shipments. And let us suppose that we are asked to predict what the frequency will be of releases of radioactivity during the shipment of spent fuel. Since we do not know this frequency for certain, we express our state of knowledge on the point by assigning a probability distribution, for example, as in Fig. 1. This distribution expresses our state of knowledge, our state of uncertainty, before we have any actual experience in shipping spent fuel. Now let us say we have 4000 shipments with no releases. How does that piece of information change our state of knowledge as expressed in Fig. 1?

To answer, we reason, very straightforwardly, as follows:

Let

B stand for the statement "we have 4000 shipments with no releases"

Let

A_1 stand for the statement "the frequency rate is 10^{-3}"

A_2 stand for the statement "the frequency rate is 10^{-4}"

A_3 stand for the statement "the frequency rate is 10^{-5}"

A_4 stand for the statement "the frequency rate is 10^{-6}"

A_5 stand for the statement "the frequency rate is 10^{-7}".

A_6 stand for the statement "the frequency rate is 10^{-8}."

With this notation, our probability distribution prior to having the evidence B can be written as

$$p(A_i) \ , \ i = 1, 2, \ldots, 6 \ .$$

After having the evidence B, our distribution becomes

$$p(A_i/B) \ , \ i = 1, \ldots, 6 \ ,$$

which we read as "probability of A_i, given B."

We seek now to find the numbers $p(A_i/B)$ in terms of the numbers $p(A_i)$. For this purpose, we note that

$$p(A_i/B)\,p(B) = p(B/A_i)\,p(A_i) \ , \qquad (4)$$

since both sides equal the probability that both A_i and B are true. Now, dividing Eq. (4) by $p(B)$, we have

$$p(A_i/B) = p(A_i)\left[\frac{p(B/A_i)}{p(B)}\right] \ , \qquad (5)$$

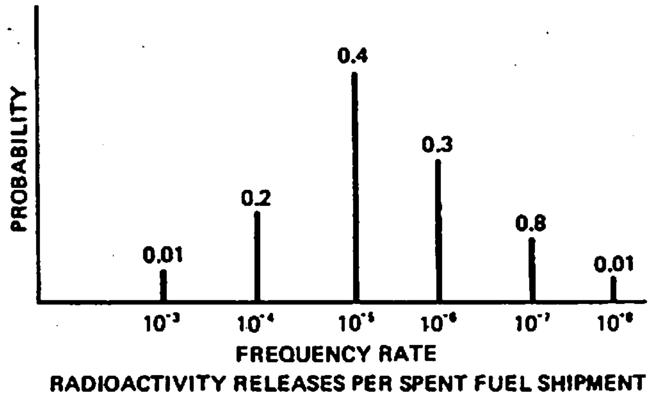

Fig. 1. Probability distribution of frequency of releases.

in which the bracketed quantity is now seen to be the correction factor that corrects $p(A_i)$ for the fact that we now have the evidence B. Equation (5), of course, is just Bayes' theorem. Now let us plug in our numbers. Consider first A_1. If the frequency of release were 10^{-3}, the probability of B, i.e., of 4000 release-free trips, would be

$$p(B/A_1) = (1 - 10^{-3})^{4000} = (0.999)^{4000} = 0.0183 \ . \qquad (6)$$

Similarly,

$$p(B/A_2) = (1 - 10^{-4})^{4000} = 0.670 \ ,$$

$$p(B/A_3) = (1 - 10^{-5})^{4000} = 0.961 \ ,$$

$$p(B/A_4) = (1 - 10^{-6})^{4000} = 0.996 \ ,$$

$$p(B/A_5) = (1 - 10^{-7})^{4000} = 0.9996 \ ,$$

$$p(B/A_6) = (1 - 10^{-8})^{4000} = 0.99996 \ .$$

Now,

$$p(B) = \sum_{i=1}^{6} p(A_i)\,p(B/A_i) = 0.907 \ ,$$

and thus our correction factors and new probabilities are given in Table I.

From this table, we can see quantitatively the effect we discussed qualitatively earlier. The introduction of information B has a small effect on our degree of confidence in the propositions A_3, A_4, A_5, and A_6, increasing these probabilities by ~10% and doing very little to distinguish among them. On the other hand, information B reduces the probability of A_2 by ~26% and virtually demolishes the probability of A_1, the correction factor in this case being 0.0202.

In words, if we have had 4000 release-free shipments, it is highly unlikely that the long-term release frequency will turn out to be in the neighborhood of 1 in 1000 shipments. Similarly, 4000 release-free shipments are enough to reduce substantially our degree of belief that the frequency is as high as 1 in 10 000 shipments. It is more likely

TABLE I

Bayesian Calculations

i	1	2	3	4	5	6
A_i	10^{-3}	10^{-4}	10^{-5}	10^{-6}	10^{-7}	10^{-8}
$p(A_i)$	0.01	0.2	0.4	0.3	0.08	0.01
$p(B/A_i)$	0.0183	0.670	0.961	0.996	0.9996	0.99996
$[p(B/A_i)/p(B)]$	0.0202	0.739	1.06	1.098	1.102	1.102
$p(A_i/B)$	0.00020	0.148	0.424	0.329	0.0882	0.01102

TABLE II

Bayesian Calculations

i	0	1	2	3	4	5	6	7	Σ
Frequency, A_i	10^{-2}	10^{-3}	10^{-4}	10^{-5}	10^{-6}	10^{-7}	10^{-8}	10^{-9}	
$p(A_i)$	0.10	0.15	0.15	0.15	0.15	0.15	0.10	0.05	
$p(B/A_i)$	3×10^{-18}	0.0183	0.670	0.961	0.996	0.9996	0.99996	0.999996	
$p(A_i)p(B/A_i)$	3×10^{-19}	0.0027	0.1005	0.144	0.149	0.150	0.100	0.050	0.696
$p(A_i/B)$	4×10^{-19}	0.0039	0.144	0.207	0.214	0.216	0.144	0.072	

in light of this experience that the release frequency will be on the order of 1 in 10^5 or less. Thus, our experience of 4000 shipments is actually very informative, and every shipment and every year that passes without accidental release will improve our confidence that spent fuel can be shipped safely without special trains.

In such applications of Bayes' theorem as this one, the questions are always asked: Where did you get the prior distribution? What if you had chosen a different one? Since the prior distribution is simply an expression of our state of certainty, we get it from within ourselves, based on the totality of our experience and judgment. It needs to come from no other place and needs no other justification. It is illuminating, however, to observe the effect of a different prior distribution. So suppose instead of the distribution of Fig. 1, we had used the distribution represented by Fig. 2, which reflects a much more uncertain state of mind than does Fig. 1. We now examine the effect of the new information on this distribution. Table II presents our new set of calculations. The new results are presented in Fig. 3.

We observe from Fig. 3 that in the 10^{-5} and above range, the distribution is different from that in Table I as a result of the different prior. For 10^{-4} and below, the distribution is the same. Thus, the story told by Bayes' theorem is the same in both cases.

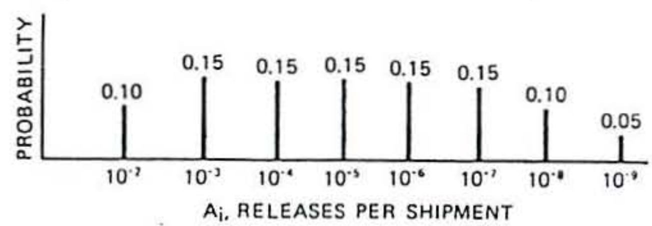

Fig. 2. Probability distribution of frequency of releases.

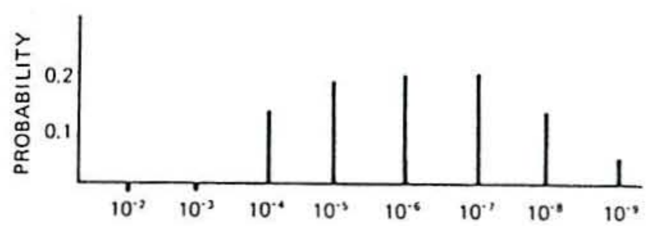

Fig. 3. Frequency of releases.

The experience of 4000 release-free shipments is not sufficient to distinguish between release frequencies of 10^{-5} and less. However, it is sufficient to substantially reduce our belief that the frequency is on the order of 10^{-4} and virtually demolish any belief that the frequency could be 10^{-3} or greater.

V. EXAMPLE 2—RELIABILITY OF A SINGLE ROW OF TURBINE BLADES

This example relates to the reliability of a single row of blades in a certain category of steam turbines for which a design improvement has been implemented. This case history shows how the reliability has improved with experience and serves as an example of how one can quantitatively assess, or predict, the reliability of equipment bases on the usually limited data in hand. It is an example of how, in developing such predictions, one needs to be careful in the interpretation of statistical data so as not to reach unwarranted conclusions.

The basic historical data for these blades are collected in Fig. 4. For each turbine, the uptimes and

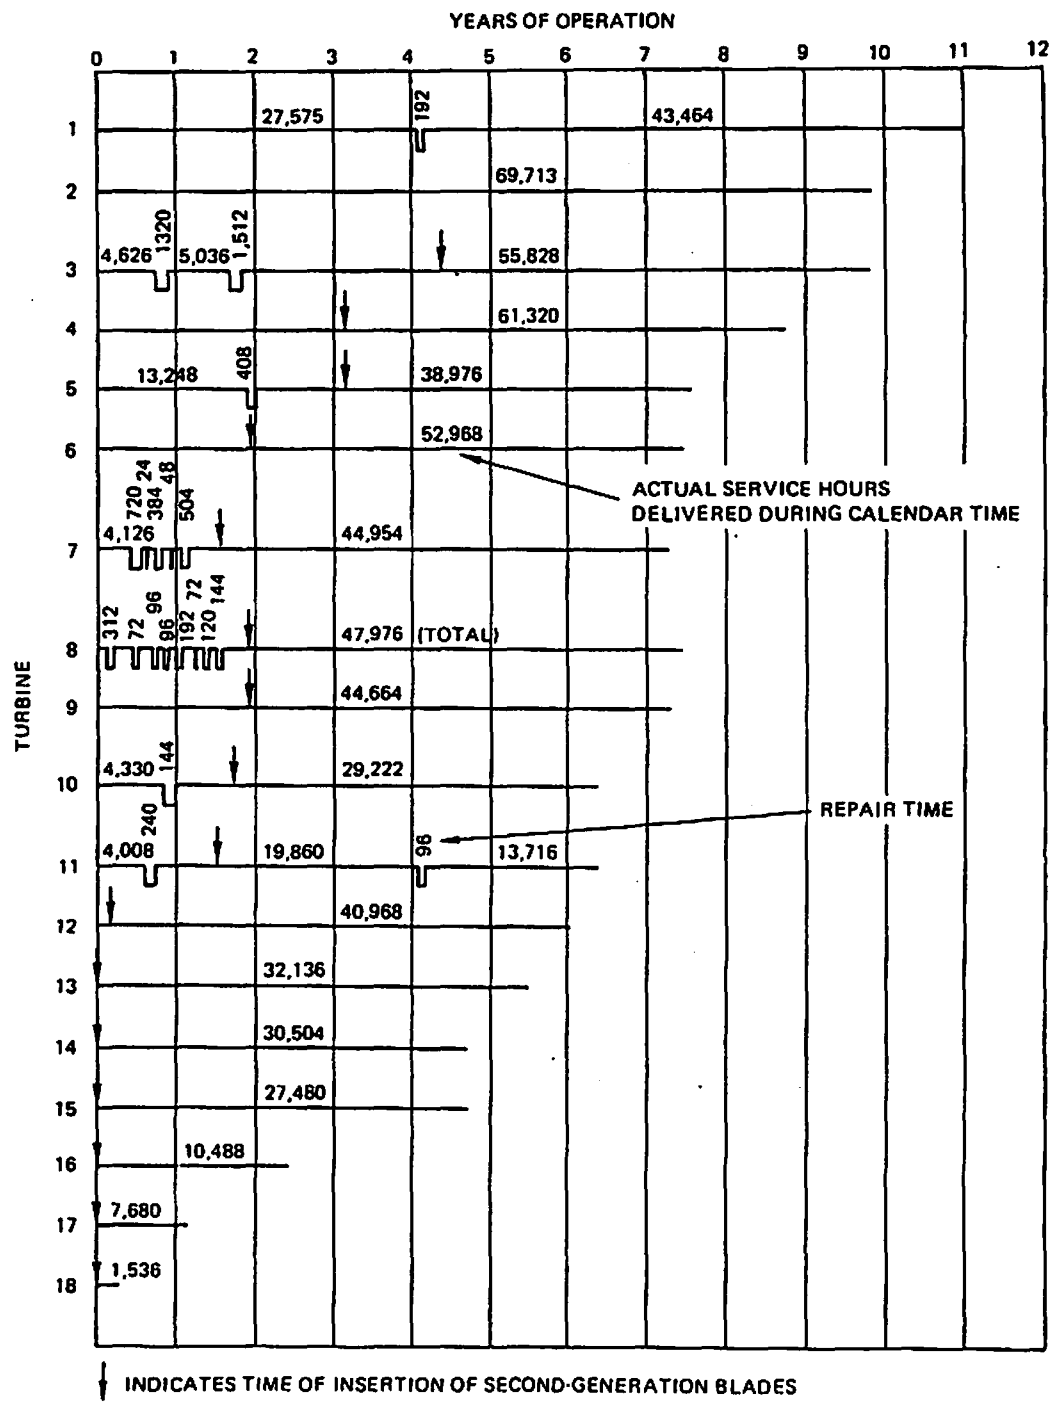

Fig. 4. Blade failure history versus calendar time.

downtimes are shown plotted against calendar years measured from the date of plant commercial operation. The numbers written on the "up" or operating parts of the times give the hours of turbine service achieved during the calendar interval between the two failures bounding the interval. The numbers listed by the "down" or forced outage intervals give the time to repair for that incident.

The 18 turbines of this type have experienced a total of 20 forced outages due to L-row blade failures; 19 of these 20 were failures of the original "first-generation" blades. The arrows in the figure designate the time at which the first-generation blades were replaced by an improved second-generation design. Note that some machines are second generation from the beginning of life, while others are entirely first generation, never having had their blades replaced.

V.A. Time to Repair (TTR)

From Fig. 4, we can abstract several interesting types of information. The first is the frequency distribution of repair times, shown in Fig. 5. From this distribution, the mean TTR (MTTR) is 330 h, with the likelihood being ~90% that the repair will take no longer than 700 h. Fifty-five percent of the time, the repair takes <200 h, and so on.

V.B. Time to Failure (TTF)

Similarly, from Fig. 4 we can abstract a TTF distribution, as shown in Fig. 6. From this distribution, the mean TTF (MTTF) is ~10 500 h. In interpreting this curve, it is important to recognize that it is based on a small amount of historical data and includes the information contained in those operating intervals that have not yet been terminated by failure. Since the curve is a frequency distribution of actual and limited historical data, not a probability distribution that reflects a state of confidence, the curve in particular should not be interpreted to mean that a failure will occur before 70 000 h. This 70 000 h is merely the largest interval observed thus far.

V.C. Availability

The preceding subsections have summarized the historical performance of this category of turbine with respect to TTR and TTF. Now what we are ultimately interested in is availability—in particular, we wish to know what will be the availability of a new turbine of this type that we might buy and install now.

To define this question more precisely, let us imagine that we do install such a machine today, and let us cast our imagination into the future, to the end of life of this machine. At that time, someone will sit down with the records and compute the average availability experienced by the machine during its life. We refer to this number as the "future lifetime availability." It is a "future-past" kind of number.

TABLE III

Unavailability Data

Turbine	Forced Outage Time (h)	In-Service Time (h)	Unavailability $U = 1 - A$
	First-Generation Blades		
1	192	71 059	0.0027
2	0	69 713	0
3	2832	26 538	0.10
4	0	23 526	0
5	408	21 120	0.020
6	0	15 066	0
7	1680	10 386	0.14
8	1014	8 184	0.11
9	0	10 980	0
10	144	9 696	0.015
11	240	8 256	0.028
12	0	1 224	0
13	0	0	---
14	0	0	---
15	0	0	---
16	0	0	---
17	0	0	---
18	0	0	---
Totals	6510	275 748	0.0236
	Second-Generation Blades		
1	0	0	---
2	0	0	---
3	0	38 952	0
4	0	37 794	0
5	0	31 104	0
6	0	37 902	0
7	0	38 694	0
8	0	39 792	0
9	0	33 684	0
10	0	23 856	0
11	96	29 328	0.0033
12	0	39 744	0
13	0	32 136	0
14	0	30 504	0
15	0	27 480	0
16	0	10 488	0
17	0	7 680	0
18	0	1 536	0
Totals	96	460 674	0.00021

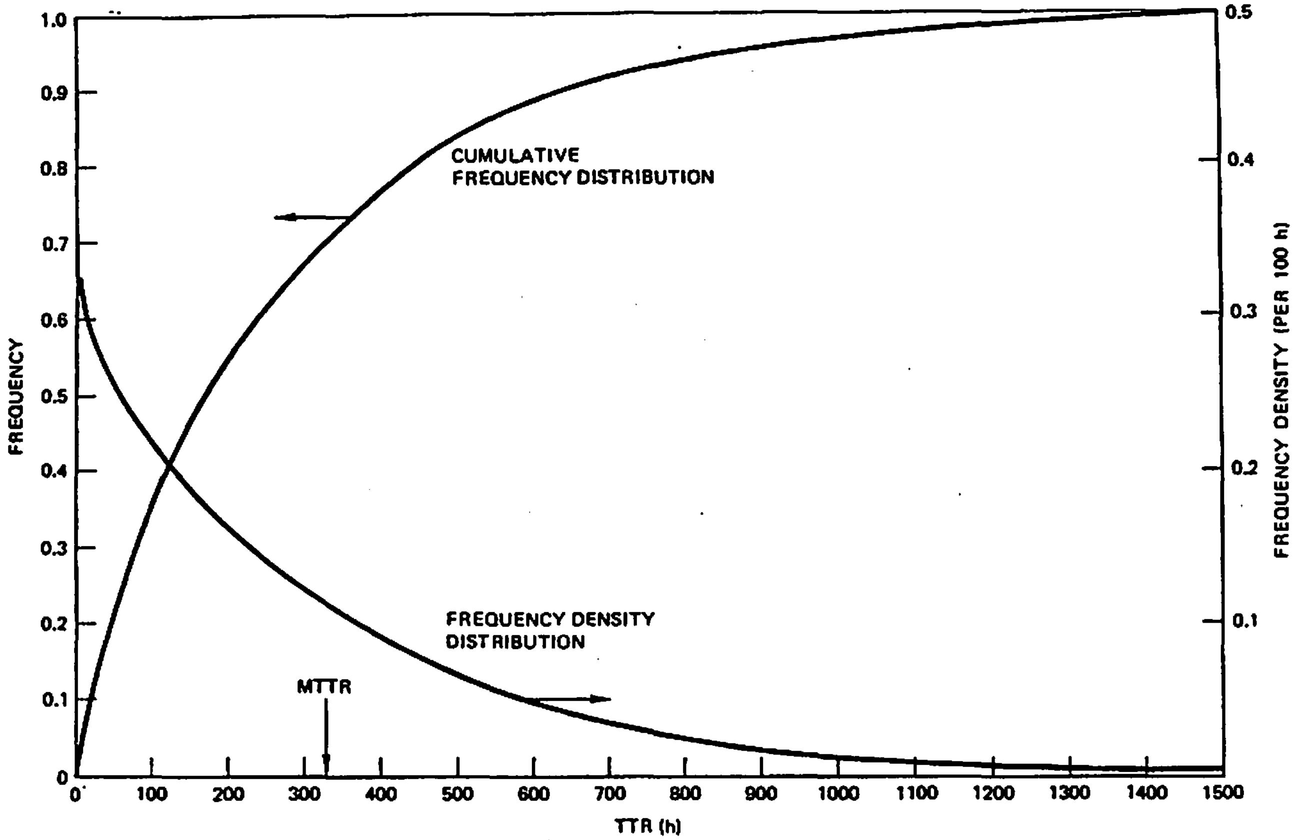

Fig. 5. Frequency distribution curves for TTR for a single row of blades.

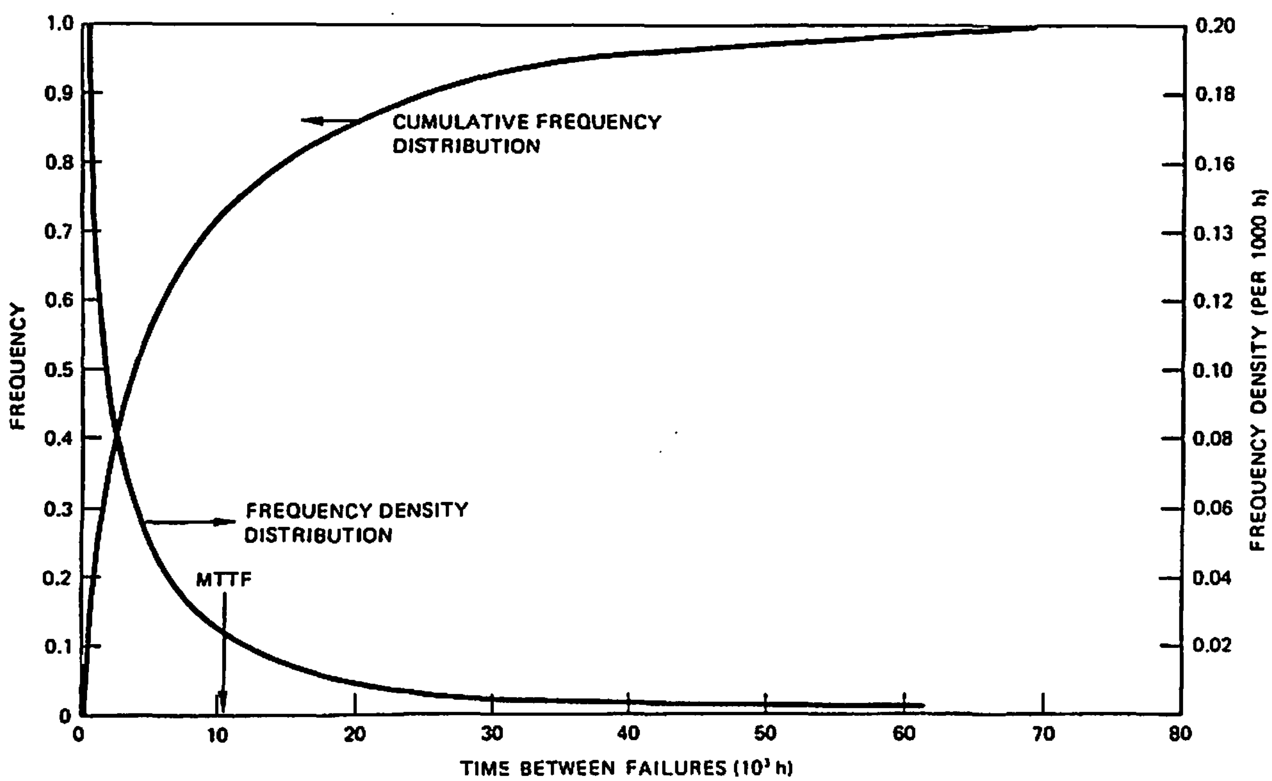

Fig. 6. Frequency distribution for time between failures.

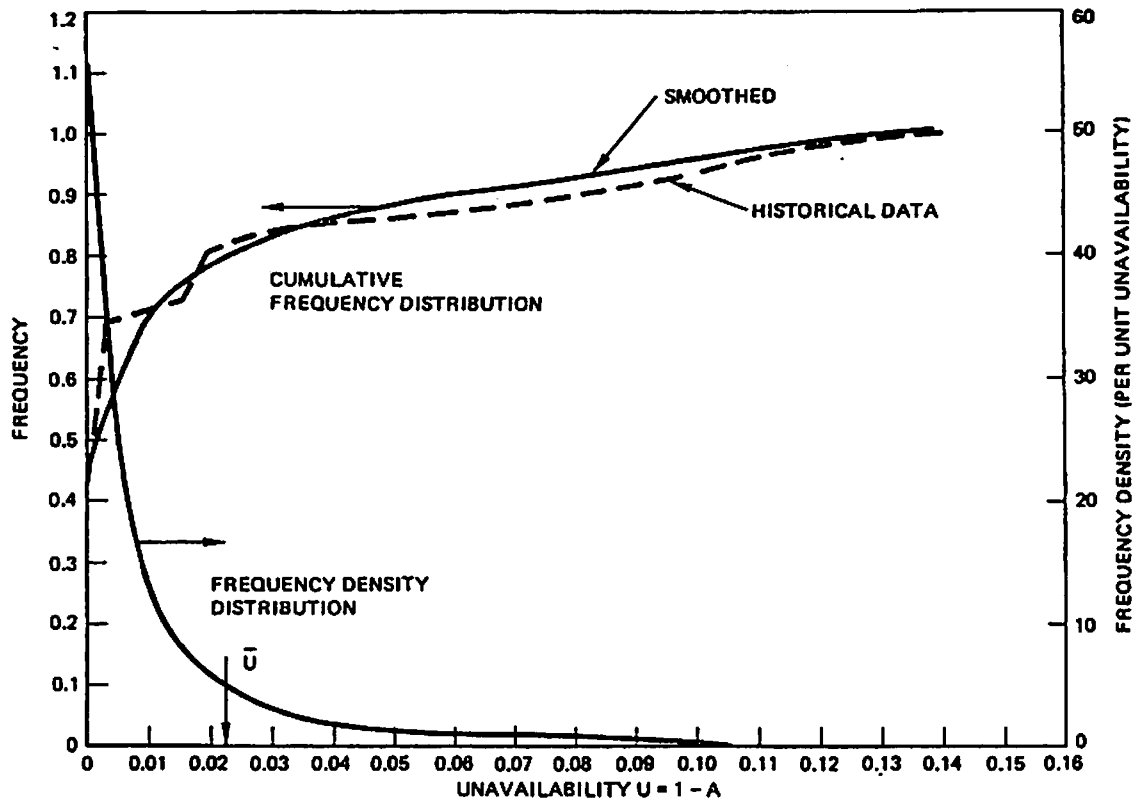

Fig. 7. Frequency distribution of unavailability: first-generation blades.

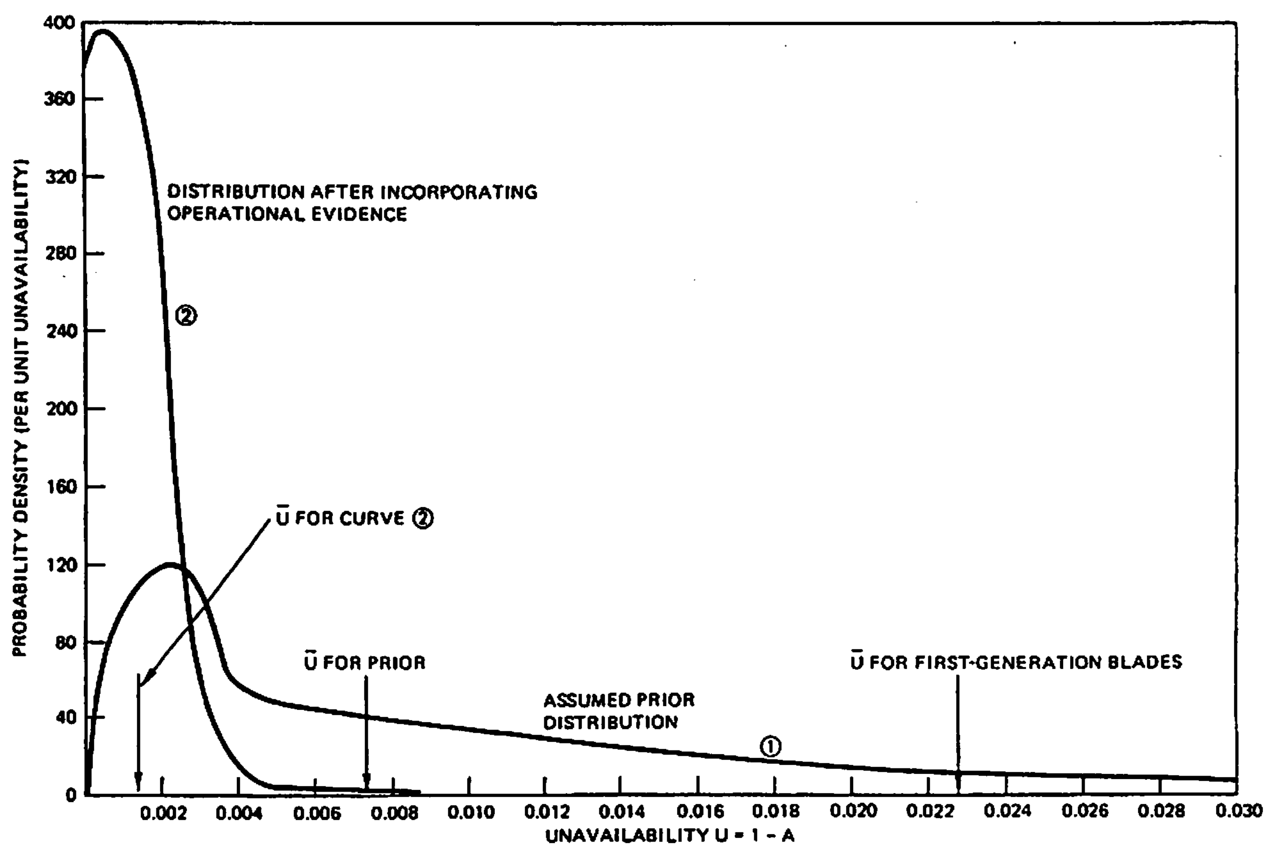

Fig. 8. Probability distribution of unavailability: second-generation blades.

Now, what we are interested in, and what we seek to do, is to predict today what that future number will be. Since this *is* a number in the future, we must, of course, express our prediction in the form of a probability distribution, and we must include in that distribution all the information, evidence, and knowledge that we have at the time we make the prediction. That is the best we can do.

There are three pieces of information we have for use in making this prediction. The first is the historical availability performance of the first-generation blades, abstracted from Fig. 4, and summarized in Table III. The second is our knowledge of the design improvements made in going to the second generation. The third, and most important, is the actual operating experience with second-generation blades. As seen in Table III, there has been 460 674 h of experience with second-generation blades, and in this time there has been only one failure. This is extremely important information, but obviously we cannot simply say that the failure rate is one per 460 674 h. We need a way of incorporating this information.

The proper conceptual tool for this purpose is Bayes' theorem, Eq. (3). To apply this theorem to our situation, we discretize the unavailability axis into definite values U_j, $j = 1, 2, 3, \ldots$. Then, letting B stand for the evidence of 460 674 h with one failure, we write

$$p(U_j/B) = p(U_j)\left[\frac{p(B/U_j)}{p(B)}\right] ,$$

where $p(U_j/B)$ is the probability we assign to the proposition that the future lifetime unavailability will be U_j after we have become aware of the evidence B. The term $p(U_j)$ is the "prior" probability that we would assign before we become aware of B; $p(U_j)$ therefore represents our state of confidence solely on the basis of the first-generation history and the design changes.

To work up to $p(U_j)$, we first plot the frequency distribution of unavailability performance of first-generation blades. This is done in Fig. 7. With this figure as background and using our knowledge of the design changes, we then judgmentally assign a prior distribution for second-generation blades. This is presented as curve ① in Fig. 8. Observe that in this curve, we express our expectation of substantial performance improvement but still in the long tail allow for the possibility that performance may not improve all that much. We shall see what happens to this tail when we incorporate the operational evidence B.

To incorporate the evidence B, we need $p(B/U_j)$, the probability that, if the lifetime unavailability were truly U_j, we would have experienced one failure in 460 674 h of operation. For this purpose, we use an exponential failure model to represent the reliability of the turbine. In this model, there is a failure rate λ_j, related to the unavailability by

$$\lambda_j = \frac{1}{\tau}\,\frac{U_j}{1 - U_j} ,$$

where τ is the MTTR. Then, if T = 460 674 h, the probability of having exactly one failure in T hours is

$$p(B/U_j) = \lambda_j T \exp(-\lambda_j T) .$$

From this relation, and noting that

$$p(B) = \sum_j p(U_j)\, p(B/U_j) ,$$

we have all we need to use Bayes' theorem.

The calculations are carried out in Table IV, and the results are plotted in Fig. 8 as curve ②. We observe in this curve that the long tail is thoroughly gone. In words, the evidence B is sufficient to demolish any belief that the unavailability would be higher than, say, 0.005. The probability of being higher than this is essentially zero. The expected unavailability, the mean of the distribution, is now

TABLE IV

Bayes' Theorem Calculations for Second-Generation Blades

(T = 460 674 h = 52.59 yr; τ = 330 h = 0.0377 yr)

j	1	2	3	4	5	6	7
U_j	0.0005	0.0015	0.0025	0.0040	0.0080	0.0160	0.0260
λ_j (yr^{-1})	0.0133	0.0398	0.0665	0.1065	0.2004	0.4039	0.6801
$P(B/U_j)$	0.3475	0.2581	0.1059	0.0207	0.0003	~0	~0
$P(U_j)$	0.08	0.10	0.12	0.16	0.24	0.20	0.10
$P(U_j)P(B/U_j)$	0.0278	0.0258	0.0127	0.0033	0.0001	~0	~0
$P(U_j/B)$	0.3989	0.3702	0.1822	0.0473	0.0014	~0	~0

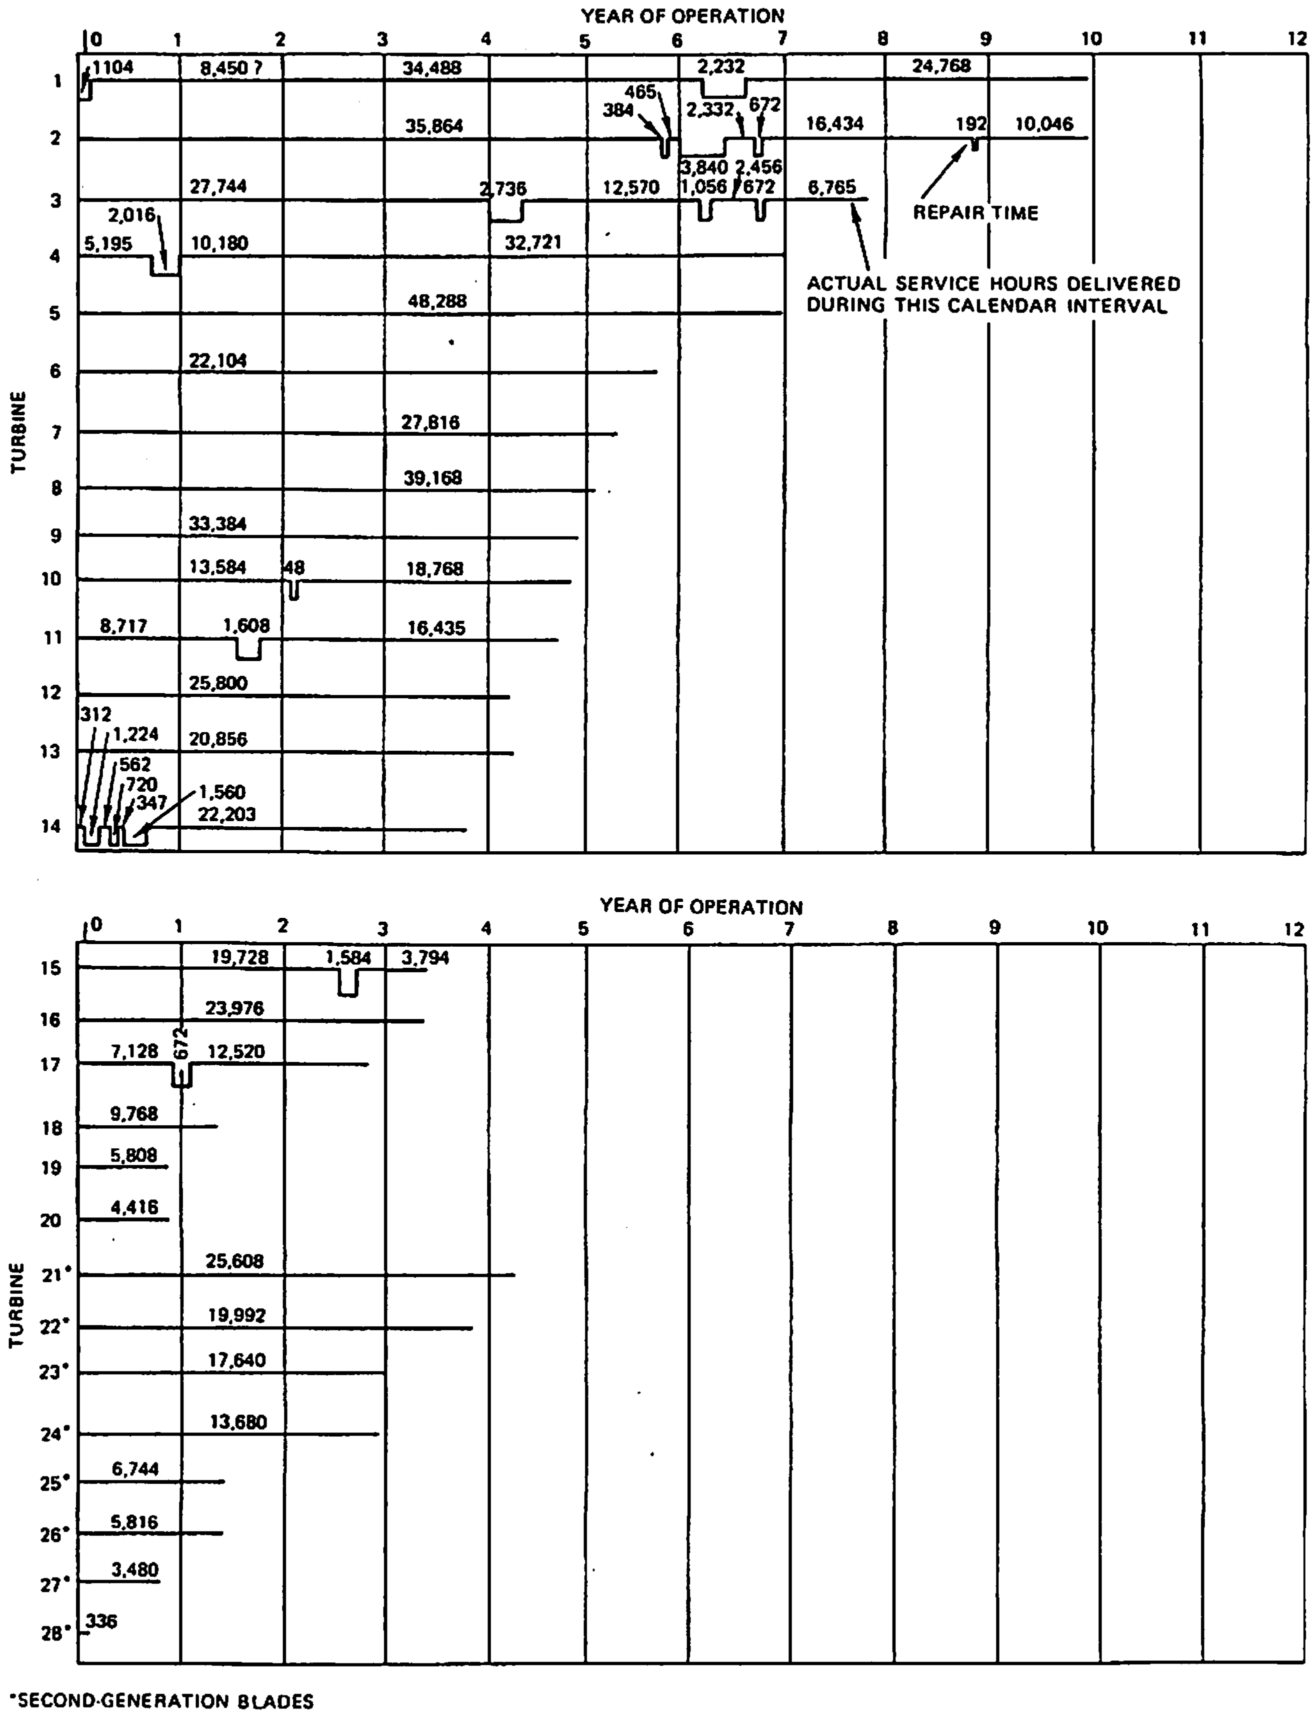

Fig. 9. Blade failure history versus calendar time.

down around 0.0014, in contrast to the expected value of 0.023 for the first-generation blades.

V.D. Recap

In this example, we have looked at historical data and inferred from the data the frequency distribution of TTR, TTF, and availability for first-generation blades. Since there was only one failure on record of second-generation blades, it was not possible to draw a frequency distribution. Instead, for this case, the inference was made using Bayes' theorem and presented as a probability distribution against lifetime unavailability. This distribution thus expresses our current state of knowledge with respect to the availability performance that will be experienced by second-generation blades. It is our prediction based on all the information available at this time. As more experience accumulates, our prediction will change.

Based on this prediction, the "expected" future lifetime availability of second-generation blades is 0.999 and the probability is essentially 100% that the availability will be 0.995 or better.

VI. EXAMPLE 3—RELIABILITY OF A CLASS OF TURBINE BLADES

This example has to do with the reliability of all rows of blades in a single style of turbine for which there are three blade design generations. For machines of the first generation, there is considerable operating experience, including a total of 17 forced outages. For second-generation machines, there are 93 296 h of service with zero forced outages. For third-generation machines, there is as yet no operating experience, although, of course, the design changes are known.

The problem is to predict, based on the information at hand, the lifetime unavailability of third-generation blades. The fundamental historical data are compiled in Fig. 9. Those turbines having second-generation blades are marked by an asterisk; the others are first generation.

Notice that the data we have cover only the first portion of life. We have no full-life-span data. Moreover, most of the data are for turbines with first-generation blades. We have only a small amount

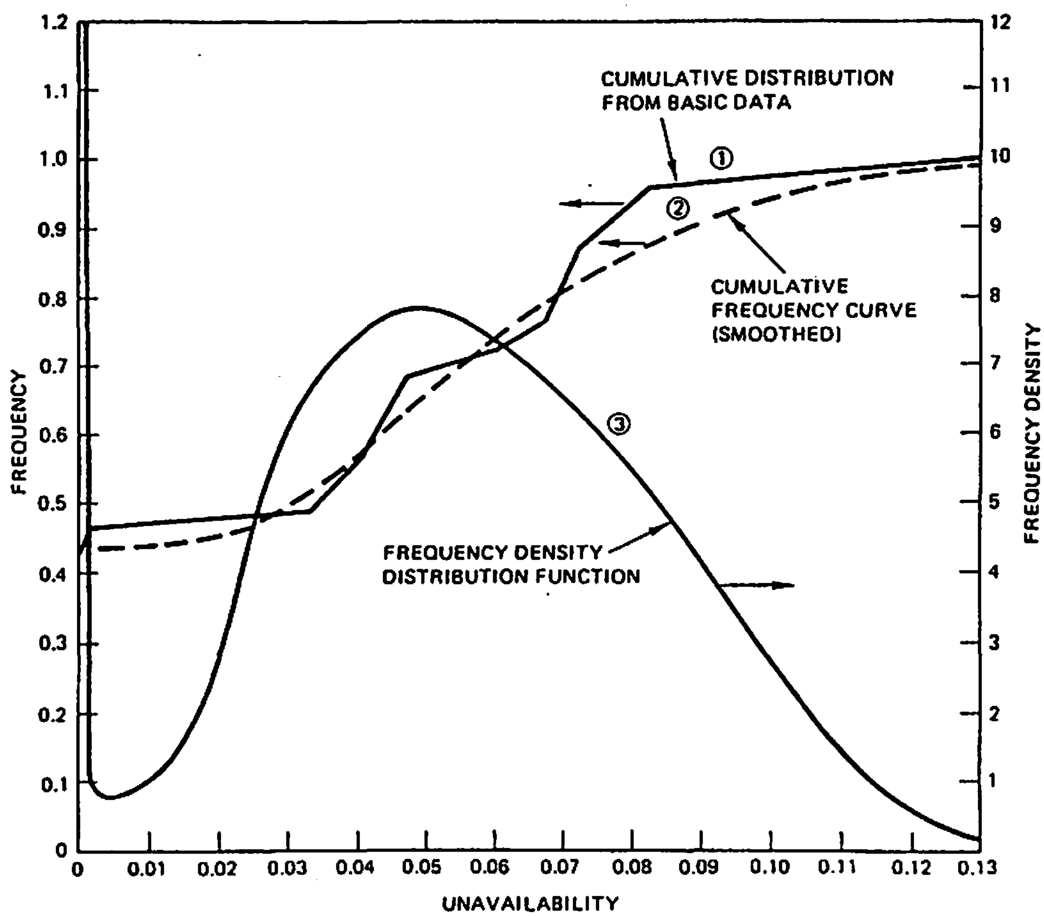

Fig. 10. Historical unavailability data: first-generation blades (using individual turbines as data points).

of data on turbines with second-generation blades, and what we are installing are turbines with third-generation blades. So, how can we predict the reliability of third-generation blades? The answer is we do the best we can, using whatever information we have and making sure that the form we choose to express our prediction also includes and expresses our total state of confidence about that prediction. We begin by calculating the frequency distribution from the first-generation data. The next step is to extrapolate from this distribution to a probability distribution for second-generation blades, solely by judgment based on the design changes made in the second generation.

Next, we incorporate the available experience with second-generation blades: 93 296 h with no failures. This is done quantitatively using Bayes' theorem. The result is an up-to-date statement of our state of certainty with respect to second-generation blades. From here, we must extrapolate to third-generation blades, again solely by judgment reflecting the presumed improvement in reliability due to the learning process, but also allowing for the possibility of some unforeseen bug, always present in a new design.

This procedure is carried out in Table V and Figs. 10, 11, and 12. Figure 10 is the historical frequency distribution of unavailability for turbines with first-generation blades. To obtain this figure, we write down the blading unavailability experienced by each turbine as in Table VI. We next consider each turbine as a separate data point and give each data point a weighting according to its total hours expressed as a fraction of the sum of the total hours over all turbines. We can then plot the cumulative distribution against unavailability, curve ①, fit this to a smoothed curve, ②, and then differentiate the smoothed curve to obtain the frequency density curve shown as curve ③.

With this curve as background, we put forth as a

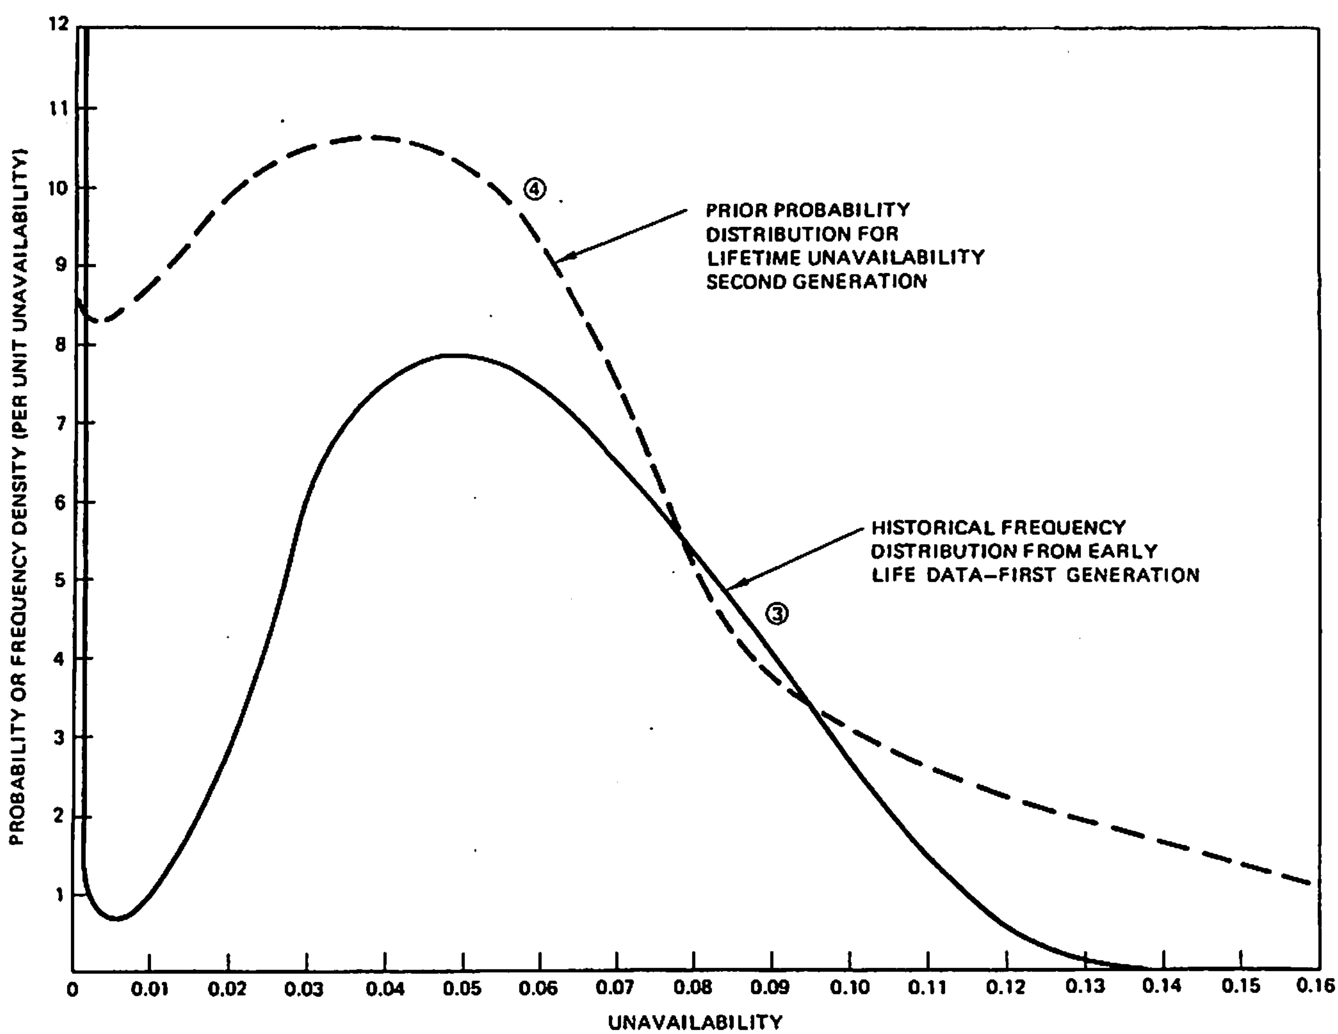

Fig. 11. Historical early life frequency distribution ③ for first-generation blading and full lifetime prior probability distribution for second-generation blading ④.

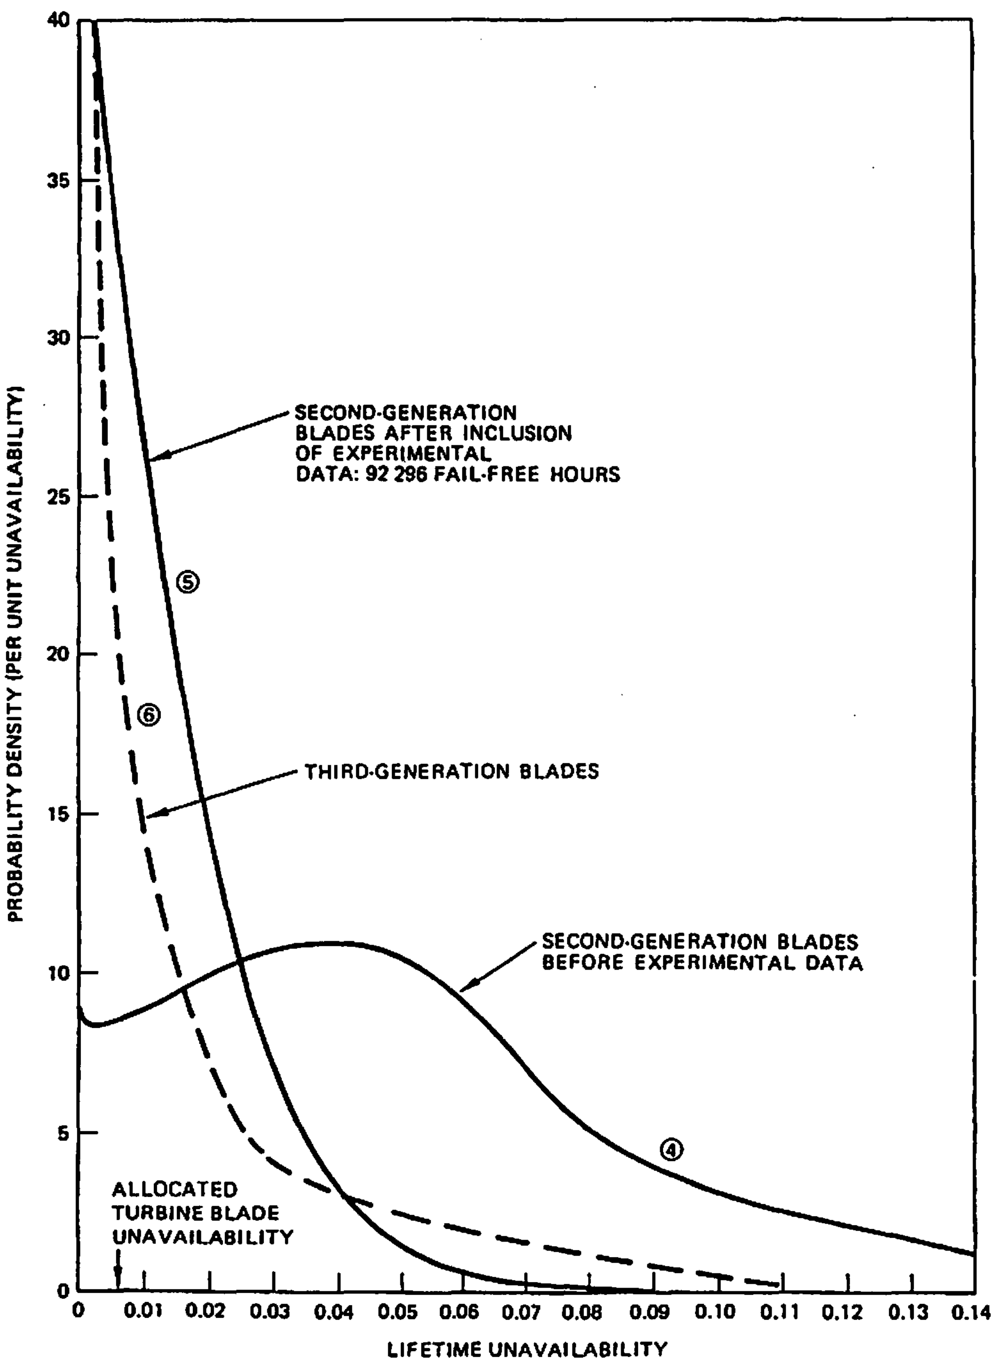

Fig. 12. Effect of experimental evidence on probability distributions for second-generation blades.

prior distribution for second-generation machines the curve labeled ④ in Fig. 11. This prior is deliberately chosen conservatively, that is, only slightly improved over the first-generation curve. The action of Bayes' theorem on this curve is calculated in Table V, wherein we now use

$$p(B/U_j) = \exp(-\lambda_j T) \ , \ T = 93\ 296\ \text{h}$$

to reflect the fact that we have zero failures in T hours.

The resulting probability distribution is displayed as curve ⑤ in Fig. 12, with an allocated or "target" unavailability of 0.0062. We observe that the inclusion of the 93 296 failure-free hours has a dramatic effect on the shape of the probability curve, shifting the bulk of the area from the 0.03 to 0.08 range to the 0 to 0.02 range. What this tells us is that while 93 296 operating hours is not yet enough to give total assurance that the allocated unavailability of 0.0062 can be met, it is certainly enough to give high

TABLE V

Bayes' Theorem Calculations for Second-Generation Blades with 93 296 Failure-Free Service Hours

i	1	2	3	4	5	6	7	8	9	10	11	
U	0	0.005	0.015	0.025	0.035	0.045	0.055	0.07	0.09	0.11	0.14	
λ_j	0	0.00280	0.00847	0.0143	0.020	0.0262	0.0324	0.0419	0.055	0.0687	0.0906	(mos^{-1})
$p(B/U_j)$	1.0	0.699	0.339	0.162	0.0760	0.0351	0.0160	0.00476	0.00089	0.000153	0.000009	
$p(U_j)$	0.01	0.085	0.095	0.103	0.106	0.105	0.10	0.148	0.072	0.055	0.091	
$p(U_j)p(B/U_j)$	0.01	0.0595	0.0322	0.0167	0.00806	0.00368	0.0016	0.00070	0.000063	~0	~0	0.1325
$p(U_j/B)$	0.075	0.449	0.243	0.126	0.0608	0.0278	0.012	0.0053	0.00047	~0	~0	

TABLE VI

Unavailability Calculations: First-Generation Blades

Turbine	Forced Outage Time (h)	In-Service Time (h)	Unavailability $U = 1 - A$
1	3 336	67 706	0.0470
2	5 088	65 141	0.0724
3	4 464	49 529	0.0827
4	2 016	48 096	0.0402
5	0	48 288	0
6	0	22 104	0
7	0	27 816	0
8	0	39 168	0
9	0	33 384	0
10	48	32 362	0.0015
11	1 608	25 152	0.0601
12	0	25 800	0
13	0	20 856	0
14	3 504	23 426	0.1301
15	1 584	23 522	0.0673
16	0	23 976	0
17	672	19 648	0.0331
18	0	9 768	0
19	0	5 808	0
20	0	4 416	0
Totals	22 320	615 966	0.0350

confidence that the unavailability of second-generation blades is <0.02, for instance, and is vastly improved over the first generation. It is seen, moreover, that this conclusion is reached in spite of the very conservative choice of prior. In fact, that choice helps to make the conclusion more evident.

Curve ⑤ is thus our current state of certainty with respect to lifetime unavailability of second-generation blades. For third-generation blades, we have no operating evidence at all at this time. Our state of confidence therefore rests mainly on improvement already achieved between first- and second-generation blades. If an equal improvement is made between second and third generations, we should meet the unavailability target handily. This optimism is tempered with the thought that the third generation is a new design and there may be unforeseen problems. We therefore express our state of knowledge with respect to turbines with third-generation blades in curve ⑥. This state of knowledge indicates that we "expect" to meet our target unavailability, although the confidence level in this expectation is only moderate at this time. Again, this state of knowledge will change as data from testing and operation become available.

VII. CONCLUSION

When questions of risk and reliability arise, they arise invariably within the context of a decision in which these factors must be traded off against cost, benefit, environment, and other factors. To make such trade-offs in a rational way and to create at least the possibility of a consensus in such decisions, it is desirable to quantify risk, reliability, and the uncertainty associated with them as completely as possible.

The appropriate framework for such quantification is the theory of probability, understood as a language for expressing states of knowledge based on information at hand. Within this theory, Bayes' theorem is extremely useful as a way of quantifying the effect of new evidence on a preexisting state of knowledge.

We have given in this paper three examples of the use of Bayes' theorem, showing how partially complete or partially relevant information can be incorporated into a full presentation of a current state of knowledge, which can then be input to a trade-off decision. We feel these examples demonstrate the effectiveness of Bayes' theorem in this regard and in making clear the exact significance of the new information. We hope, therefore, that these examples will contribute to furthering the cause of rationality and explicitness in (especially public) decision-making.

REFERENCES

1. T. C. FRY, *Probability and Its Engineering Uses*, 2nd ed., p. 128, D. Van Nostrand Company, New York (1964).

2. *Van Nostrand's Scientific Encyclopedia*, 2nd ed., p. 154, D. Van Nostrand Company, New York (1967).