U. S. AIR FORCE

PROJECT RAND

RESEARCH MEMORANDUM

APPLICATION OF THE BAYES' TECHNIQUE
TO SPARE-PARTS DEMAND PREDICTION

W. H. McGlothlin
Eloise E. Bean

RM-2701

January 3, 1961

Assigned to ______________________________

This research is sponsored by the United States Air Force under contract No. AF 49(638)-700 monitored by the Directorate of Development Planning, Deputy Chief of Staff, Development, Hq USAF.

This is a working paper. It may be expanded, modified, or withdrawn at any time. The views, conclusions, and recommendations expressed herein do not necessarily reflect the official views or policies of the United States Air Force.

The RAND Corporation
1700 MAIN ST. • SANTA MONICA • CALIFORNIA

47

SUMMARY

This Research Memorandum describes a practical method for predicting spare-parts demand rates, which may be readily carried out with a few simple, manual computations. The method is an application of Bayes' Theorem of classical probability theory, which combines a priori information and data derived from subsequent events in order to evaluate the likelihood of future events. A previous Research Memorandum,* to which this is a supplement and companion-piece, provided the mathematical explication of the Bayes' Theorem and its applicability to demand prediction. The present study was written in response to the need for a simpler, non-technical, more directly usable version, containing more illustrative material.

The method takes advantage both of initial estimates and of unfolding operational experience -- even when this experience would be dismissed as too "limited" under current Air Force practices. It should therefore be particularly useful during the early life of a weapon system, when demand data are meager but fairly large commitments must nevertheless be made for spare parts.

The current Air Force procedure normally uses the initial estimate of demand made at the provisioning conference until the sample of operational experience is judged "extensive" -- representative of the true demand rate. This may result in continuing use of the initial estimate even after its information value has come to be overshadowed by that available from demand experience. The Bayesian technique, on the other hand, produces a weighted average of the initial estimate and that computed from subsequent demand

*W. H. McGlothlin and R. Radner, The Use of Bayesian Techniques for Predicting Spare-Parts Demand, The RAND Corporation, Research Memorandum RM-2536, March 1, 1960.

experience. The weights depend on the expected accuracy of the initial estimate in comparison with the accuracy of the rate computed from a given sample of operational experience; thus there is a gradual shift from the initial estimate to the computed rate, based on the estimated relative value of the two pieces of information. In cases when the early demand experience is considered nonrepresentative for reasons other than the size of the sample, it would still be disregarded or modified. In particular, a limited use of the procedure is explained for the case when age-dependent demand processes prevent the extrapolation of early demand rates to later periods.

The values of the parameters needed for the procedure are based on empirical measures from several weapon systems. The expected accuracy of the initial estimate is based on a comparison of actual estimates with the demand rates which developed after relatively large amounts of operating experience. An empirical evaluation of demand variability provides a measure of the expected accuracy of demand rates computed from small samples of operating experience.

The study also discusses the use of a confidence interval computed about the estimate, which takes into account the uncertainty of both the initial estimate and the demand rate computed from the sample. Thus we may choose to use an upper confidence limit as the demand-rate estimate, and make an approximate probability statement concerning the likelihood that the "true" rate exceeds the estimate.

The study gives the detailed steps of the Bayesian procedure and includes the graphs needed for applying it to aircraft parts. The method can also be used for missiles or other equipment, but the graphs may require modification when operation is measured in units other than flying-hours.

CONTENTS

I. INTRODUCTION

BACKGROUND OF THE STUDY

This study is a simplification of and supplement to its companion-piece, RM-2536.* RM-2536 provided a mathematical exposition of the Bayes' Theorem of classical probability theory, specifically as applied to spare-parts demand prediction. The object of this Research Memorandum is to provide procedures and statistics for direct use in demand prediction.

The RM has two principal differences from its companion-piece: (1) its explanation of the rationale underlying the Bayes' Theorem and its application is largely non-technical; (2) as an instruction manual might do, it outlines a detailed step-by-step procedure for applying the technique. For convenience, this Research Memorandum repeats some parts of its predecessor, but the reader is referred to the latter for a mathematical development of the subject.

The study thus joins a series of RAND publications on demand prediction, the number of which reveals the continuing importance both RAND and the Air Force attach to the subject.** Previous studies have examined the feasibility of predicting demand from past operational experience, and to some extent from *a priori* estimates made at a provisioning conference. The contribution of the present paper is a prediction method based on the integration of the *a priori* estimate with the rate computed from the operational experience.

---

*See footnote, p. iii.

**For example, see Bernice B. Brown, *Characteristics of Demand for Aircraft Spare Parts*, The RAND Corporation, Report R-292 (ASTIA No. AD 107426), July, 1956; T. A. Goldman, *Relationships Between Program Elements and System Demand for Airframe Spare Parts*, The RAND Corporation, Research Memorandum RM-1858 (ASTIA No. AD 123514), January 22, 1957; *idem*, *A Priori Demand Prediction: A Case Study of B-52 Airframe Parts*, The RAND Corporation, Research Memorandum RM-2088 (ASTIA No. AD 144299), January 10, 1958; *idem*, *System Demand for Airframe Parts During Weapon-System Phase-in* (U), The RAND Corporation, Research Memorandum RM-2381, May 18, 1959 (Secret); and J. W. Petersen, *Savings from Procurement Deferral with Interim Contractor Support: The Case of High Value Airframe Spares*, The RAND Corporation, Research Memorandum RM-2085 (ASTIA No. AD 144297), January 10, 1958.

APPROACH OF THE STUDY

This paper examines a specific aspect of the supply problem: estimating spare-parts demand rates when the data from operational experience are still relatively sparse, and reliance must continue on the demand-rate estimate made by the initial provisioning conference. The current Air Force practice is to use the provisioning-conference estimate as long as operational experience is "limited," and switch to the computed demand rate only when operational experience is judged "extensive." (The computed demand rate is simply the number of observed demands divided by equipment operating time.) The dynamic procedure recommended in this paper provides for a systematic shift from the initial estimate to the computed demand rate as operational experience increases. It does so by employing a weighted average of the initial estimate and the computed demand rate. At any particular time, the weights are intended to reflect the usefulness of the two pieces of information for predicting the future demand rate.

For example, suppose the initial estimate of the demand rate made at the provisioning conference for Part X is 1.0 -- one unit per 1000 hours of operation -- but the observed demand during the first 5000 hours is two, giving a computed rate of 0.4. What should we predict the future demand rate to be? The current Air Force procedure would probably judge the demand experience to be limited and merely retain the 1.0 demand rate. The Bayesian procedure would base its prediction on a weighted average of 1.0 and 0.4. The weights would depend on three factors: (1) the estimated accuracy of the initial estimate, (2) the inherent variability of demand from one period to the next, and (3) the size of the operational sample. This study will provide the practical procedures for determining these weights and the resulting estimated demand rate. They are intended primarily for use with the more

expensive parts (over $200); first, because of the major savings possible in this category, and second, because the assumptions underlying the methods are more nearly met for these parts.

Section II of this paper considers demand processes, first as independent of part age and then as a function of part age, and discusses the implications of each situation from the standpoint of measurement and prediction of demand. Section III provides a non-mathematical description of the Bayesian technique for predicting demand rates. Section IV presents empirical measures of the parameters needed to develop the weights used in the procedure. Section V provides the graphs, formulae, and the step-by-step procedure for making demand-rate predictions. A formula for determining the statistical confidence of these predictions is given, along with suggestions for employing confidence intervals in demand predictions. Examples are provided for three types of application of the Bayesian procedures. Section VI offers some conclusions on the usefulness of the procedure.

II. DEMAND PROCESSES

BASING PREDICTIONS ON PAST EXPERIENCE

Before we turn to the Bayesian procedures, it will be well to discuss an important assumption underlying demand prediction: Whenever we use operational experience from a past period to predict spare-parts demand rates for a future period, we accept the assumption that the processes producing the demands remain substantially the same in both periods. This assumption is the source of much of the difficulty in demand prediction. If we accept the assumption when it is not true, our predictions are likely to be grossly incorrect. If we dismiss the assumption as false when actually it is true, we deprive ourselves of highly useful information for making future estimates.

Errors of the second type can cause gross over-procurement of spare parts in the Air Force. Highly inflated estimates of initial demand rates may be used far into the operational program because of a vague assumption that demand is positively correlated with part age, and that future estimates therefore cannot be based on past experience. In actual fact, the demand rates for many spare parts are virtually independent of part age, at least [redacted] for the typical, rapidly-obsolescing weapon systems.* A recently advocated method of dealing with age- or wear-dependent demand is discussed in Sec. V, together with a way of articulating it with the Bayesian technique.

[redacted] *See, for example, Milton Kamins, Determining Checkout Intervals for Systems Subject to Random Failures, The RAND Corporation, Research Memorandum RM-2578, June 15, 1960. Kamins points out that, historically, it has been found that many components, sub-systems, and systems experience exponential failures, which have been described as "random," "accidental," "chance," and "Poisson."

CONVENIENT CATEGORIES OF DEMAND PROCESSES

For the purpose of discussing _when_ we may assume that future demands will be positively related to past demands, it is convenient to categorize demand processes as follows:

I. Probability of demand is independent of part age and time of observation.

II. Probability of demand is independent of part age and time of observation, after adjustment for short-interval scheduled maintenance (2 to 12 months).

III. Probability of demand is not independent of part age, because of initial heterogeneity of the part sample.

IV. Probability of demand is not independent of time of observation, because of changes in the external environment.

V. Probability of demand is not independent of part age, because of internal changes in the part, or relatively long intervals (greater than 12 months) between scheduled maintenance.

Case I is the simplest from the standpoint of demand measurement and prediction. The demand data[*] are converted either to a demand rate per operational unit or to the inverse, mean-time-to-demand (issue interval). If a part has survived t hours, the probability that it will survive to $t + \Delta t$ hours is the same, whether $t = 1$ or 1000 hours.

In Case II, malfunctions are again independent of part age, but their correction is often delayed until normal scheduled-maintenance periods. No adjustment is necessary if the ratio of hours of operation to number of equipments undergoing scheduled maintenance during the period is approximately equal to the interval between maintenance. It typically is, after

[*]As used here, the term "demand" will exclude non-recurring demands such as those "... caused by initial activation, increases in allowances, special projects, the assembly of technical order compliance (TOC) kits, and the prestocking of mobilization reserves." The quote is taken from AMCM 57-1, Hdqs. AMC, Wright-Patterson Air Force Base, Ohio, p.73.

the equipment has been in operation for some time. An adjustment often is required during the phase-in period, however, upon which this paper focuses. (An adjustment formula is suggested in Sec. V.) But again, if the interval between scheduled maintenance is no more than one or two months, the unadjusted demand rate is usually adequate.

In Case III, the parts in the initial sample have heterogeneous resistance to operational stress. Parts with low resistance tend to fail sooner and more often than do the stronger parts. The surviving group thus emerges with a higher average resistance to stress than the original sample had, and the demand rate will therefore decrease as a function of part age. This is known as a "debugging" or "burn-in" phenomenon; it is generally important only in early periods of operation. The demand data may be adjusted by excluding both the operational hours and the resulting demands during the burn-in period. The demand rate for subsequent periods would then be considered independent of part age. Without this adjustment, the error in future demand predictions would be on the high and thus conservative side.

In Case IV, changes in demand rate are related to environmental changes. For example, improvements in operation or maintenance would probably decrease demand rates. Similarly, many technical changes are usually made in military equipment after it becomes operational, and they often affect the demand rate for individual parts. Such changes cannot generally be related to part age, since they affect new and old parts alike. We normally resort to a moving-average measure of demand in this situation, and use only the more recent data (usually from the past 12 months) to compute the demand rate.

Finally, in Case V, changes in demand rate are related to wear or aging in parts. To predict demand rates in the first three cases, the sample

demand data could be used directly or after minor adjustments. In Case IV, if we predict future demand on the basis of the sample data we would normally over-estimate demand, which at least would be an error of the conservative type. In Case V, the sample demand data collected when parts are of one age do not necessarily indicate what demand rate we should expect when the parts are older.

There are two ways to predict demands in Case V. One is to determine the demand rate as a function of part age from the operational data, and then predict future demands from this empirically developed function, taking into account the expected age of the parts in the prediction period. The other method is to make an *a priori* assumption about the age/demand relationship and then predict future demand in the same way. The advantage of the first method is that it provides a means of measuring age-effects in the demand process. It has two disadvantages. First, both the data collection and statistical computations for time-referenced failure data are expensive in comparison to the normal procedure of obtaining only total demands and operating time. Second, it usually takes so long to arrive at a satisfactory measure of the age/demand function that it often has limited use in provisioning for modern, rapidly obsolescing weapons. For many representative cases, about half the parts in the sample must fail before enough data are available to develop an adequate age/demand function. Of course, information gained about a certain part by this technique would be useful in predictions for similar parts in future systems.

In summary, if we propose to use early phase-in demand data to revise initial estimates of demand rates, we must assume a demand process that is independent of part age, at least during the age intervals for which the

[redacted]

prediction is made. If the demands arise mainly from long-interval scheduled maintenance, or if the part is expected to have a wear-out pattern on the basis of previous experience with similar parts, then there is no alternative to introducing the age factor in the demand-prediction process. Many parts, however, particularly in the very important electronic area, have demand patterns which are essentially independent of part age.[*] In general, it seems expedient to assume that demand is independent of age when there is no objective evidence to the contrary, because it is so advantageous to base revised estimates on as much early demand data as possible.

---

[*] Aeronautical Radio, Inc., Inter-Base Report No. 1 -- A Preliminary Study of Equipment Reliability, March 15, 1955. Cf. also Kamins, op. cit., pp. 1-3.

III. PREDICTING DEMAND RATES FROM INITIAL ESTIMATES AND DEMAND DATA

Suppose the initial provisioning conference for a particular type of aircraft has predicted that the demand rate for Part X will be μ_0. Further, for the sake of simplicity, suppose that these aircraft have flown a total of 1000 hours a month for the first five months of operation. We have observed the number of demands during this period and can compute the demand rate per 1000 hours, $\bar{D}$. Now we would like to re-examine our prediction of the demand rate for Part X, and perhaps make a revised estimate, μ_1. Several approaches are possible. As mentioned earlier, the current Air Force procedure would probably ignore the computed demand rate, $\bar{D}$, and continue to use the initial estimate, μ_0, for procurement purposes on the basis that 5000 hours of operation constitutes "limited" experience. Another approach might be to "split the difference," arriving at the revised estimate, μ_1, by adding μ_0 and $\bar{D}$ and dividing by two -- a simple arithmetic average. The approach advocated in this paper would go one step further and make the revised estimate, μ_1, a weighted average of μ_0 and $\bar{D}$, in which the respective weights, w_0 and w_s, would reflect the accuracy of μ_0 and $\bar{D}$. Thus we may write:

$$\mu_1 = w_0 \mu_0 + w_s \bar{D} \quad ,$$

where the weights w_0 and w_s add to 1; i.e.,

$$w_0 + w_s = 1.$$

The current Air Force procedure of retaining the initial estimate when the operational experience is judged "limited" could be written:

$$\mu_1 = 1.0\ \mu_0 + 0.0\ \bar{D} \doteq \mu_0$$

and when operational experience is judged extensive:

$$\mu_1 = 0.0\ \mu_0 + 1.0\ \bar{D} = \bar{D}.$$

Similarly if we "split the difference " between μ_0 and $\bar{D}$ we have:

$$\mu_1 = 0.5\ \mu_0 + 0.5\ \bar{D}\ .$$

The following discussion will show how the weights, w_0 and w_s, depend on the relative accuracies of the initial estimate, μ_0, and of the computed demand rate obtained from a given amount of operational experience. It is intended only for explanatory purposes, and should not be confused with the empirical procedure for determining the weights described in Sec. IV.

First, suppose we are able to determine the accuracy of an initial estimate of the demand rate, by examining the difference (error) between the estimate and the "true" demand rate. Of course we never know exactly what the "true" rate is, but after obtaining the demand experience for a very large number of operating-hours, we are able to get a close estimate of it. For a large sample of parts, we could plot a frequency distribution of these differences such as that shown in the left half of Fig. 1. According to this hypothetical error distribution, we would expect 20 per cent of our initial estimates to be within ±1 unit of the "true" demand rate, 55 per cent to be within ±3 units, etc.*

Next we wish to determine the accuracy of a demand rate computed from a

*The total area in the error distribution represents 100 per cent of the cases.

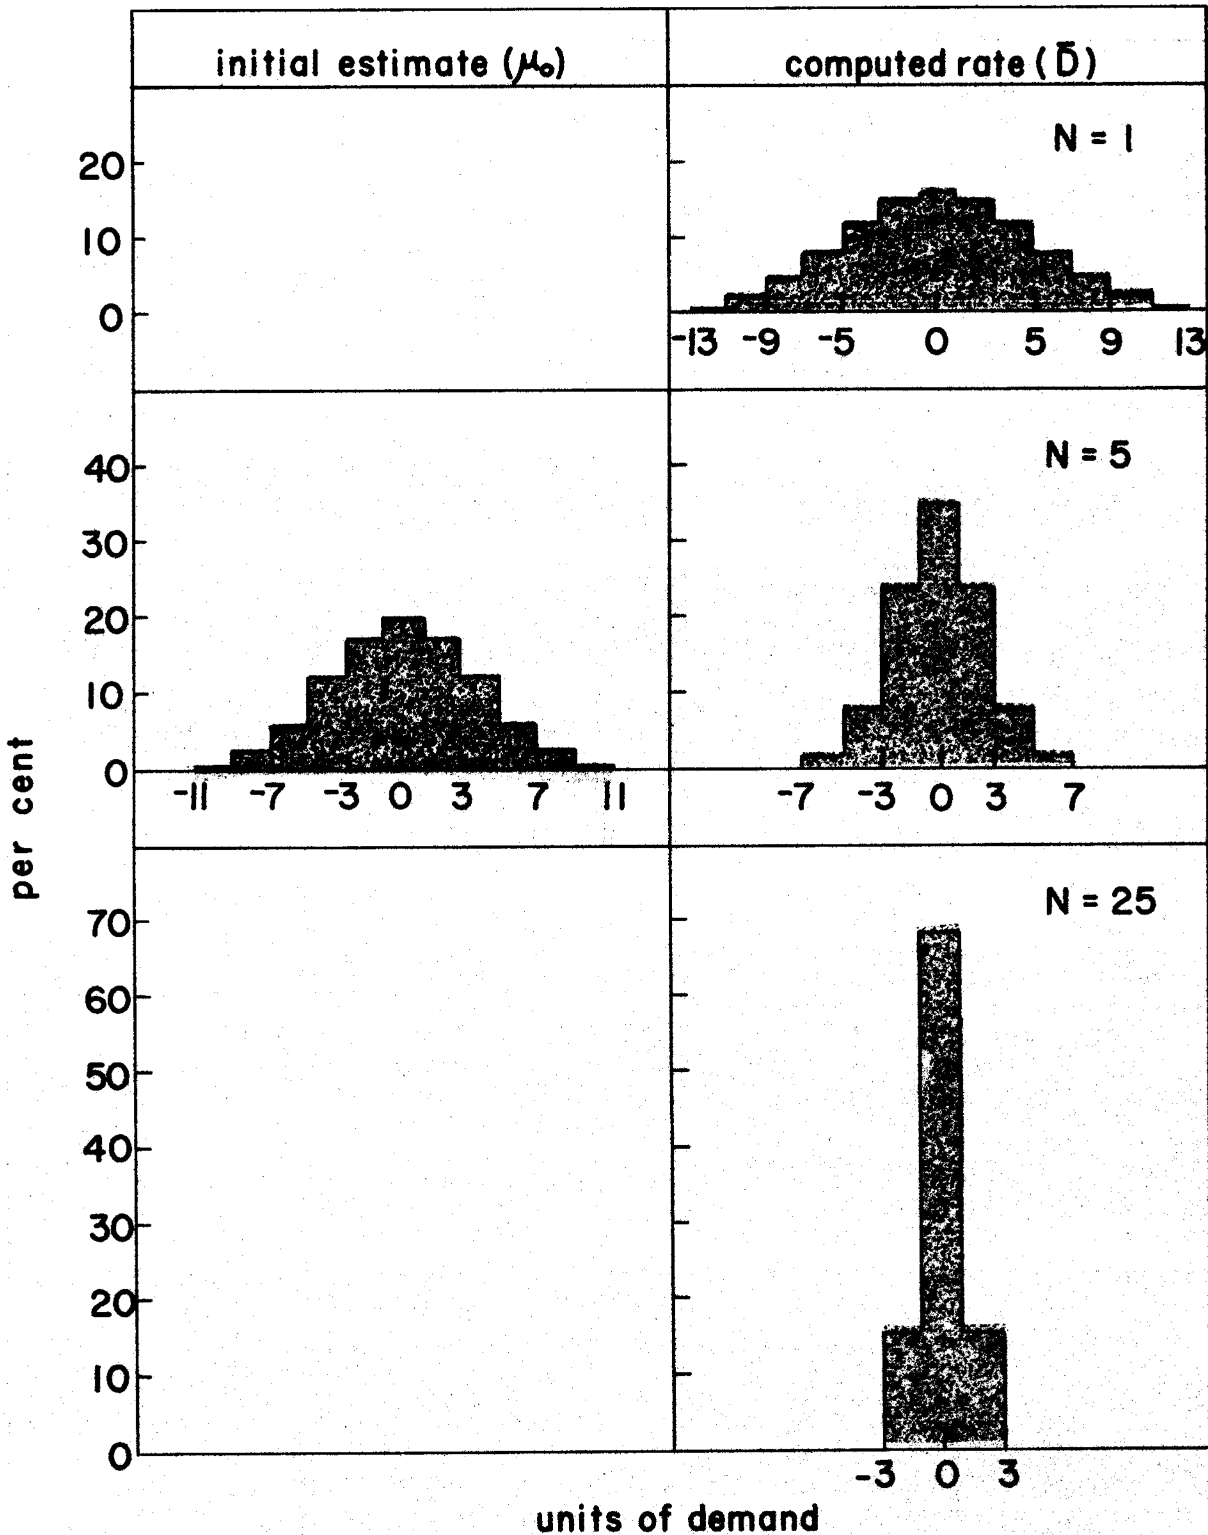

Fig. 1—Hypothetical frequency distribution of errors for initial estimates and computed demand rates

sample of operational experience, $\bar{D}$. Again, we take the demand rate for a very large number of operating-hours as the "true" rate. Suppose we then divide our large sample of operating-hours into 1000-hour periods. We could then observe the number of demands that occurred in the first period, second, etc., and obtain the differences between the number of demands occurring in the various 1000-hour periods and the "true" rate per 1000 hours as computed from the total sample. When we plot these differences as a frequency distribution we have a description of the accuracy of a demand-rate prediction based on the demand observed for a single 1000-hour period drawn at random. A hypothetical error distribution of this type is shown in the upper right portion of Fig. 1. We would expect the demand observed for a single 1000-hour period (N=1) to be within ± 3 units of the "true" rate 45 per cent of the time. Note that the error distribution for $\bar{D}$ depends on how widely the observed demand fluctuates from one 1000-hour period to the next. With large fluctuations, the demand observed for a single period will be a poorer estimate of the "true" rate than if the demand were more uniform.

Now suppose that we observe the demand during five 1000-hour periods (N=5). Our $\bar{D}$ value would be the total observed demand divided by 5. Intuitively we would feel that the $\bar{D}$ based on an N of 5 is likely to be a better estimate of the "true" rate than is the demand for N=1. Actually, the error distribution of $\bar{D}$ will decrease at the rate of $1/\sqrt{N}$.* Note the $\bar{D}$ error distribution for N=5 and N=25 in Fig. 1.

We now see that the error distribution for $\bar{D}$ depends on (1) how widely

*The reader who has had a beginning course in statistics will recognize that the concept we are discussing here is the relation of the standard error of the mean to the standard deviation of the demand distribution, i.e.,

$$\sigma_m = \frac{\sigma}{\sqrt{N}}$$

the month-to-month demand fluctuates, which is described by the error distribution for N=1, and (2) the amount of operational experience (N) upon which $\bar{D}$ is based. The $\bar{D}$ error distributions for N=5 and 25 in Fig. 1 are completely determined by our beginning distribution for N=1.

Suppose we determined a $\bar{D}$ error distribution for every N between 1 and 25. We could then take the μ_0 error distribution and find a $\bar{D}$ error distribution that was approximately the same. Since μ_0 and $\bar{D}$ are both estimates of the "true" demand rate, if they have the same error distributions they [redacted] must be equally good estimators. Suppose the μ_0 error distribution is the same as that for $\bar{D}$ when N=3. The predictive value of μ_0 would then be worth 3000 hours of operational experience. If we think of the worth of both μ_0 and $\bar{D}$ in terms of equivalent values of N, we can now determine what weights (w_0 and w_s) we should assign to μ_0 and $\bar{D}$.

Suppose that after 3000 hours of operational experience we compute a $\bar{D}$. We also have a μ_0 estimate which we assumed is worth 3000 hours of experience. The weights to be attached to μ_0 and $\bar{D}$ are then:

$$w_0 = \frac{3}{3 + 3} = 0.5 \qquad w_s = \frac{3}{3 + 3} = 0.5 \quad .$$

Our best estimate of the "true" demand rate, μ_1, is then:

$$\mu_1 = 0.5\,\mu_0 + 0.5\,\bar{D}$$

Suppose that instead of 3000 hours experience we have 6000 hours. This is equivalent to having two $\bar{D}$ estimates, each based on 3000 hours. Our μ_0 estimate is still equivalent to only 3000 hours. Our weights would now be:

$$w_0 = \frac{3}{3 + 3 + 3} = 0.33, \quad w_s = \frac{3 + 3}{3 + 3 + 3} = 0.67 \quad .$$

We may now write general equations for w_0 and w_s. Let R equal the number of operational hours (expressed in thousands) to which the predictive value of μ_0 is equivalent, and N equal the actual amount of operational experience upon which $\bar{D}$ is based. Then:

$$w_0 = \frac{R}{R + N} \quad , \quad w_s = \frac{N}{R + N} \quad .$$

The value of R in practice will depend on (1) the error distribution we assume for μ_0, and (2) the error distribution for $\bar{D}$ when N=1. Recall that the $\bar{D}$ error distributions for $N \gg 1$ are determined by that for N=1, and the latter is simply representative of the demand variability from one 1000-hour period to the next.

Figure 2 shows w_s plotted as a function of N for four values of R (remember that $w_0 = 1 - w_s$). In the following Section we shall examine the R values determined from actual initial estimates and demand experience.

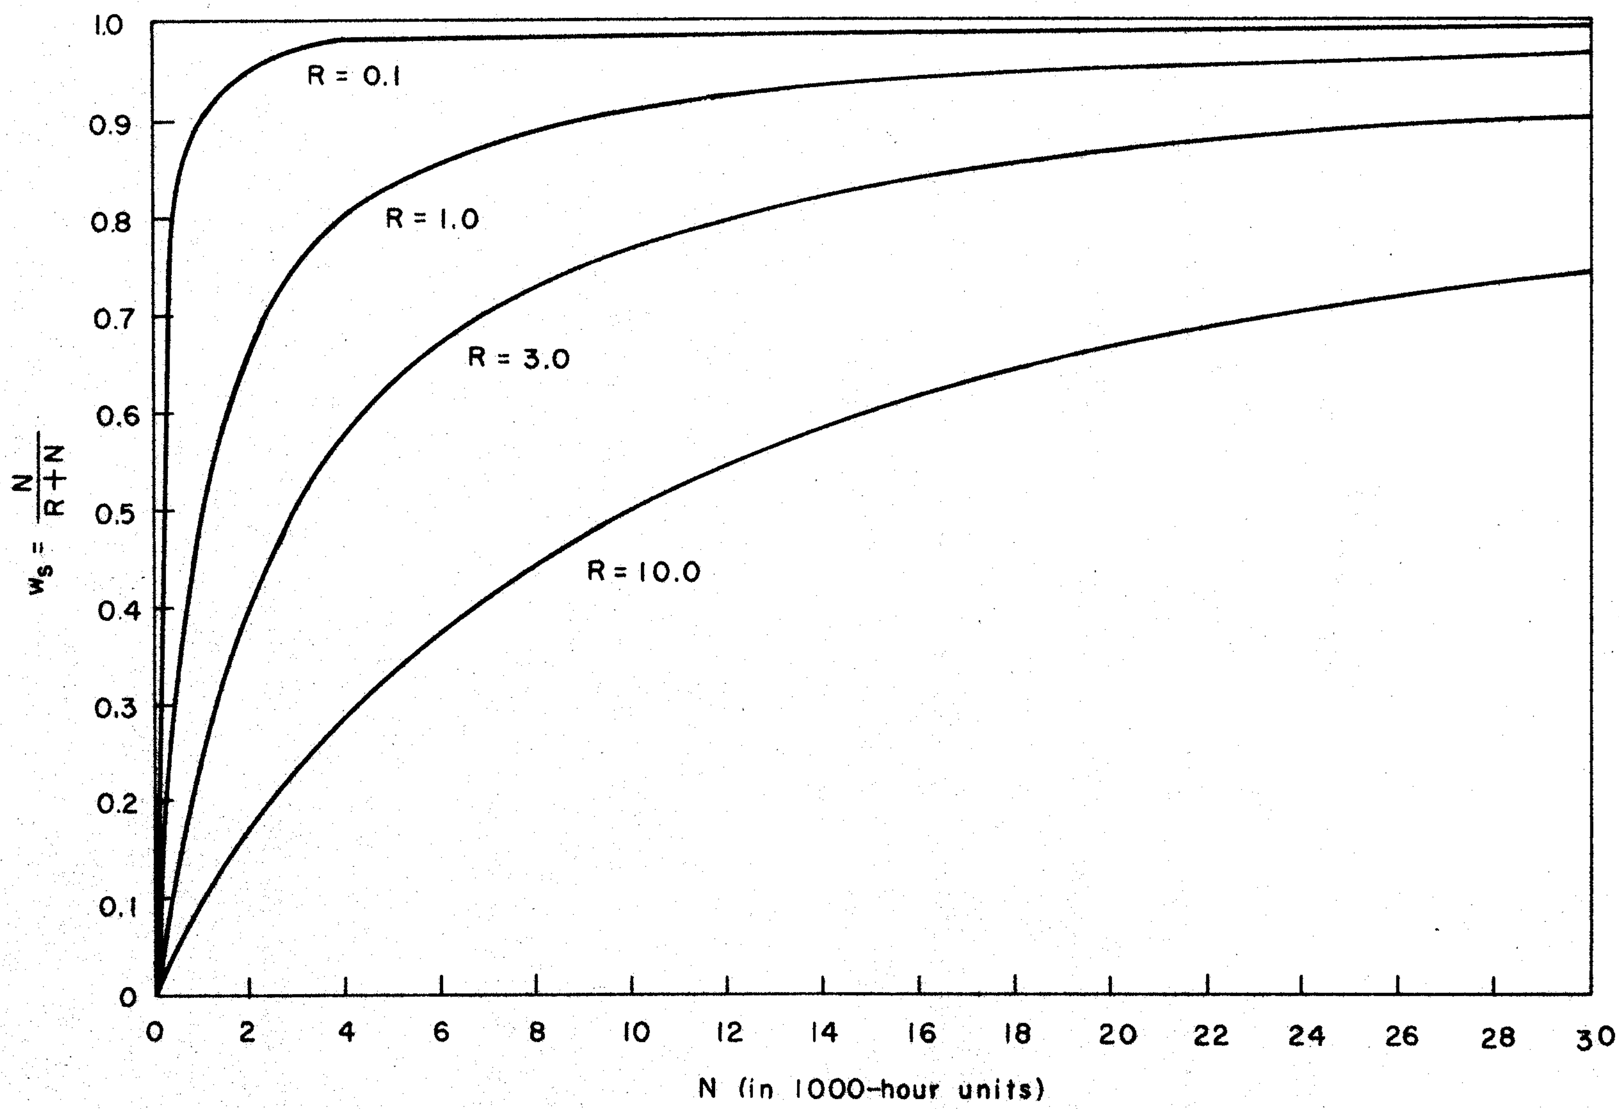

Fig. 2 — w_s as a function of N

IV. EMPIRICAL MEASUREMENT OF THE PARAMETERS

PARAMETERS FOR THE BAYESIAN PROCEDURE

In order to use the Bayesian procedure to estimate demand rates, we require four parameters:

(1) μ_0 = initial estimate of the demand rate (from the provisioning conference);

(2) $\bar{D}$ = demand rate computed from the operational data;

(3) N = the amount of operational experience expressed in 1000's of flying-hours or similar units;

(4) R = the amount of operational data which is equivalent in predictive value to that of the initial estimate.

The first three parameters are immediately available. The fourth is the crucial part of the Bayesian procedure in that it, in conjunction with the value of N, permits us to assign relative weights to μ_0 and $\bar{D}$ which reflect their respective information value for predicting the "true" demand rate. It is not statistically feasible to determine R by comparing empirical error distributions of μ_0 and $\bar{D}$ as was done for explanatory purposes in Sec. III. Rather, we shall examine empirical data to obtain estimates of these error distributions, and from these get our desired value of R.

ERROR DISTRIBUTION ABOUT $\bar{D}$

Recall from the example in Sec. III that the error distribution about $\bar{D}$ for N=1 is simply a frequency distribution of the differences between the observed demands for the 1000-hour periods and the "true" demand rate per 1000 hours. One descriptive statistic of this error distribution is the average of these differences. This would provide us with an average expected error when we use the demand observed for an N=1 to estimate the "true"

demand rate. For reasons beyond the scope of this discussion, a more useful measure of the $\bar{D}$ error distribution is an average of the square of the differences rather than the average difference. We shall call this measure s^2.

Ideally, we should like to have an estimate of s^2 for the particular part for which the demand rate is being estimated. However, the prediction technique described here is useful only when we have an initial estimate and a relatively small amount of subsequent operational experience. Under such conditions, we do not have sufficient data to compute an s^2 value for the part in question. From the analysis of other parts, however, we shall find that s^2 can be estimated as a function of the demand rate. We have an estimate of the demand rate, of course, so by determining a general relationship between s^2 and the demand rate, we may estimate the value of s^2 for the particular part.

The relation between s^2 and the demand rate was examined for two sets of data. The first was made up of 104 B-52 parts, drawn from four property classes (Airframe, Fire Control-Gunnery, Communications, and Accessories).* The second was composed of 24 components from an air-to-air missile. The B-52 data were obtained from a special base-level demand-reporting system covering 34 months, in which 69,000 hours were flown. The missile data resulted from approximately 30,000 checkouts made over a period of 26 months.

Figure 3 shows s (rather than s^2) plotted as a function of the demand rate per 1000 flying-hours, $\bar{D}$, for the 104 B-52 parts.** The function fitted

---

*These data were generously made available from another study being conducted by Bernice Brown and James Houghten in the Supply Group of RAND's Logistics Department.

**For the method of computing s (the standard deviation of the demand distribution), see McGlothlin and Radner, RM-2536, pp. 28-35. In Figs. 3 and 4, s rather than s^2 is plotted as a function of $\bar{D}$ because the resulting function is more nearly linear.

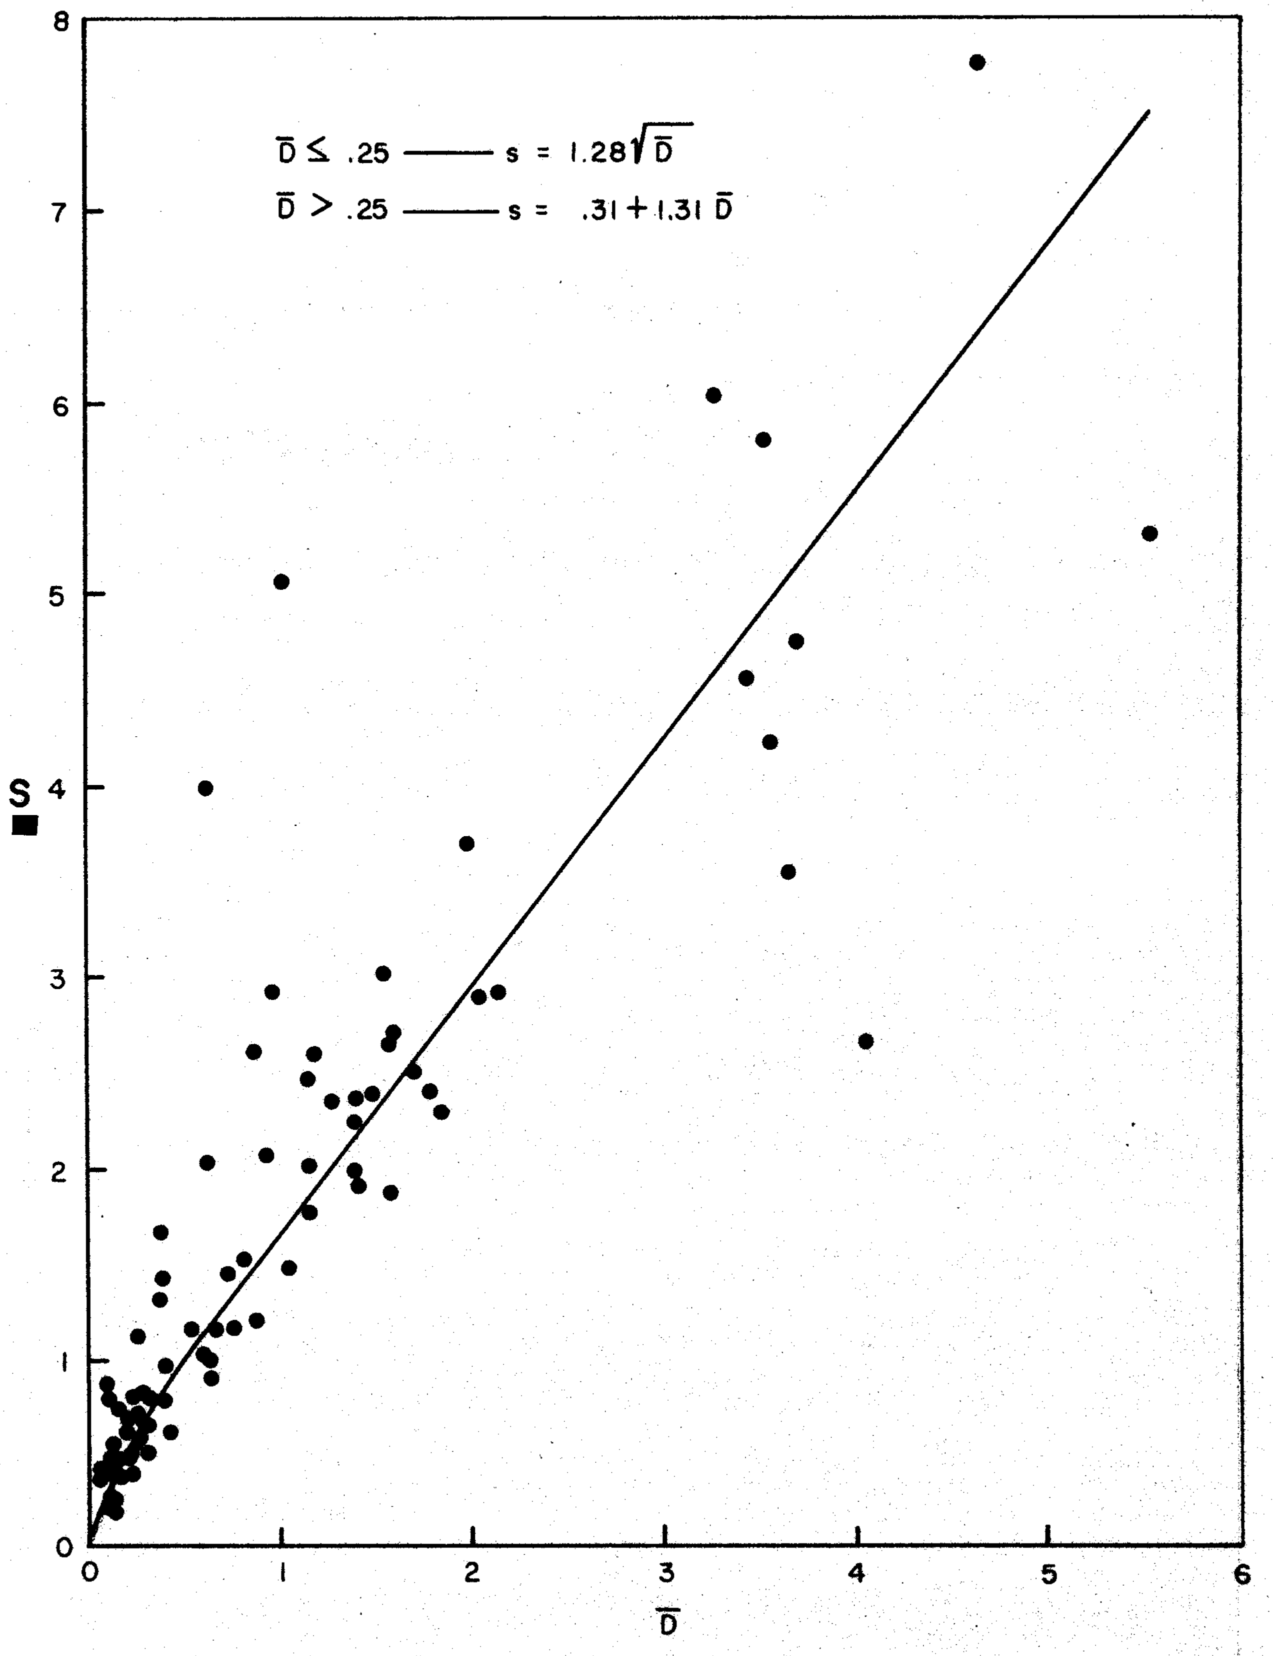

Fig. 3—Standard deviation (s) as a function of the mean demand ($\bar{D}$) for 104 B-52 parts

to the data is in two parts. For $\bar{D} \leq 0.25$, $s = 1.28\sqrt{\bar{D}}$; and for $\bar{D} > 0.25$, $s = 0.31 + 1.31\,\bar{D}$. The reader should note that while the term $\bar{D}$ represents the computed demand rate as before, it is now based on a very large N and [redacted] hence is a close estimate of the "true" demand rate. The $\bar{D}$ value which will normally be available for the Bayesian weighting procedure will be based on a much smaller N and hence will be a poorer estimate of the "true" rate.

We believe that the relationship between s and $\bar{D}$ shown in Fig. 3 is fairly characteristic of aircraft spare parts.* In any event, our ultimate aim is to estimate R, which is dependent on the error distributions for both μ_0 and $\bar{D}$. As will be seen in the following Subsection, our present ability to evaluate the former is quite poor, so attempts to gain additional accuracy in the estimate of s are probably not warranted.

Figure 4 shows s plotted as a function of the demand rate per 200 checkouts for the 24 missile parts. It will be noted that the same function that was fitted to the B-52 data also describes the missile data fairly well. It should be remarked, however, that had we chosen a unit of N other than 200 checkouts, the slope of the empirical function fitted to the data would have been different. This points up a difficulty in the use of the empirical parameters provided in this Section for equipment other than aircraft. If the Bayesian procedure is being used for demand-rate prediction for aircraft parts, then the operational unit used should be 1000's of flying-hours, so that the graphs provided in this Section can be used directly. If another operational unit is to be used such as missile-alert-hours, it will probably be necessary to recompute s as a function of $\bar{D}$ for the data in question.

---

*The authors found the same type of relationship between s and $\bar{D}$ in an unpublished analysis of 50 C-47 parts over a period of 36 months.

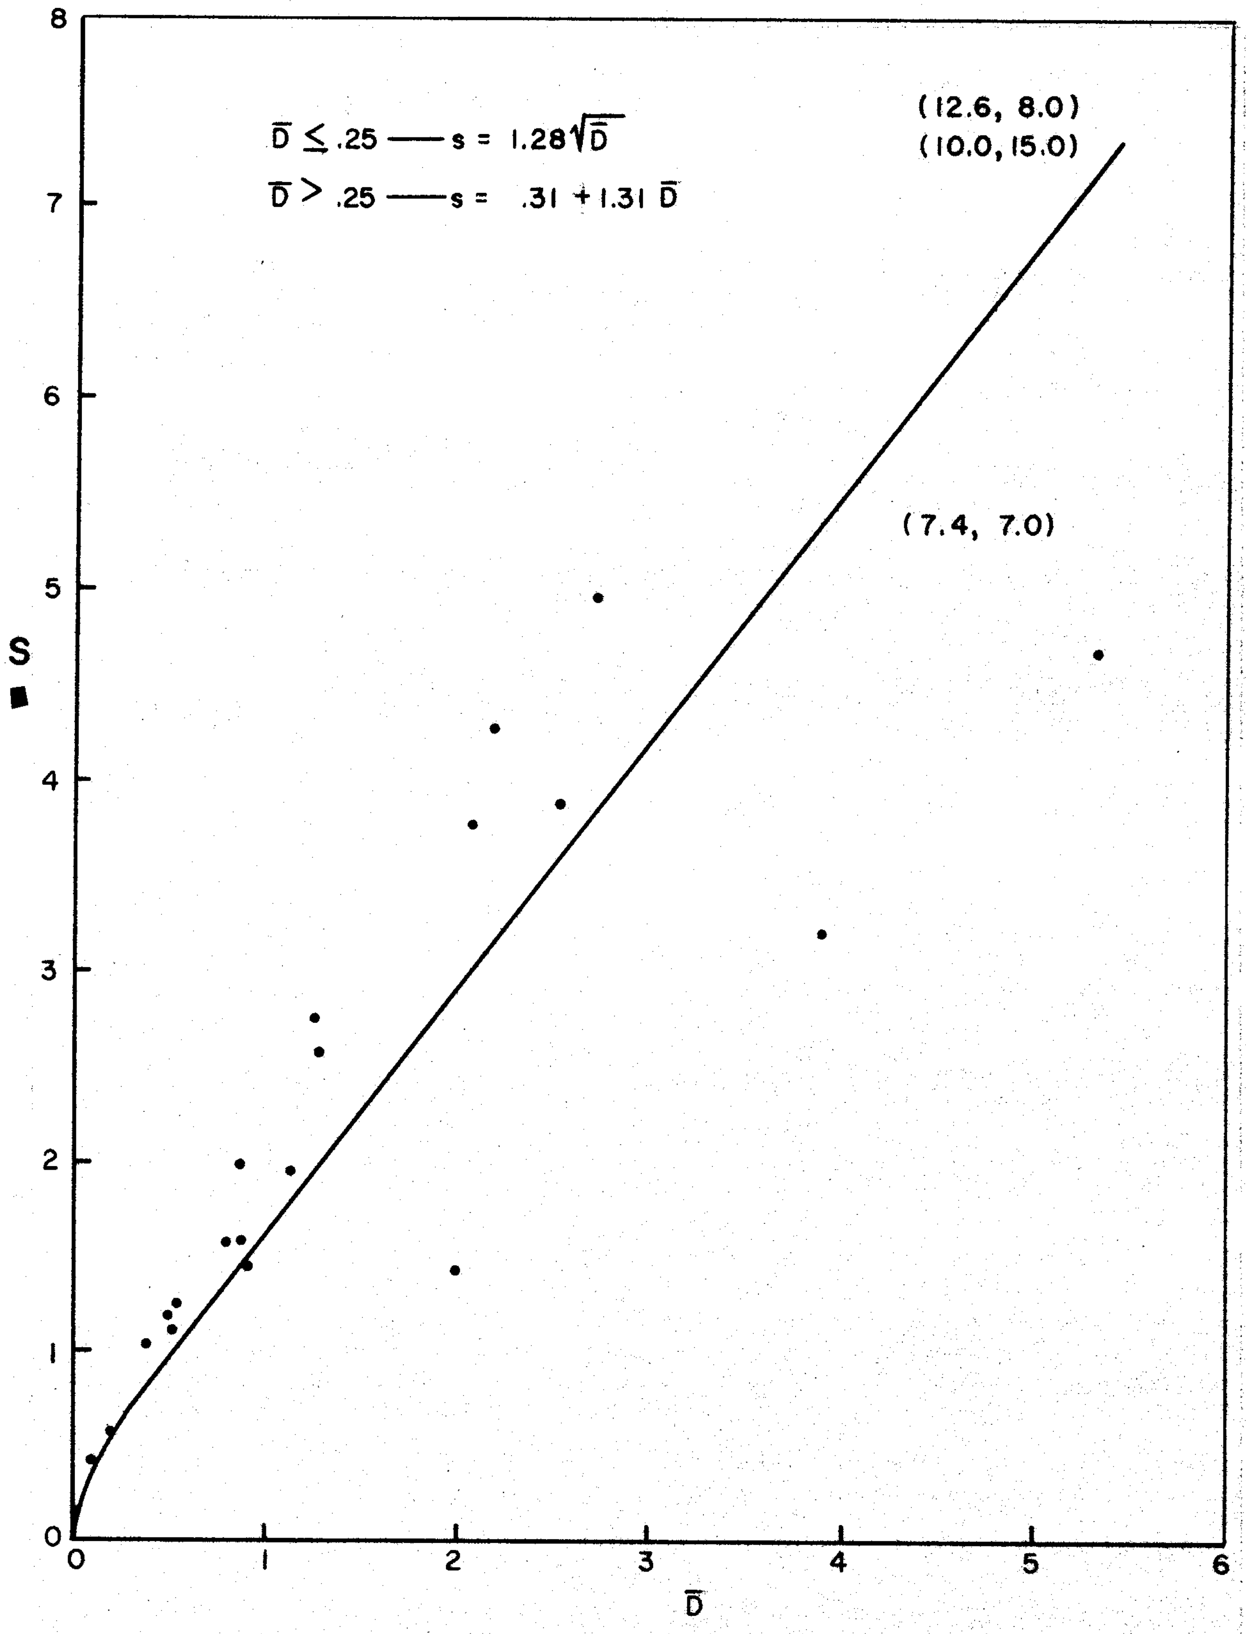

Fig. 4—Standard deviation (s) as a function of the mean demand ($\bar{D}$) for 24 missile parts

ERROR DISTRIBUTION ABOUT μ_0

In order to obtain an empirical measure of the error distribution about μ_0, we must determine, for a group of parts, the differences between the initial estimates and the computed rates which develop after a large number of operating-hours. We shall assume that the computed rates are good estimates of the "true" demand rates. Again, we will measure the average of the square of the differences rather than the average of the differences. We will call this measure σ_0^2; it is analogous to the s^2 values obtained for the $\bar{D}$ error distribution.

We should like to have the spare parts divided into at least two groups: one for which the estimator has relatively high confidence in his estimate, and a second group with which he has had less experience and thus has lower confidence in the accuracy of his estimate. It would appear reasonable that the estimator could make such a crude differentiation.* At present, however, the Air Force initial provisioning conference does not attach any such confidence ratings to initial estimates, so we will have to examine the parts as a single group.

Two sets of data were available: the initial demand-rate estimates for 61 of the previously described sample of B-52 parts,** and the initially estimated and computed demand rates for 46 B-66 and C-133 Hi-Valu parts.*** For the B-52 parts, Fig. 5 presents a scatter diagram of initial estimates versus the demand rates computed from 69,000 flying-hours.**** Three types

---

*Goldman, RM-2088.

**The authors wish to thank C. Tokarz and H. Allen for their cooperation in obtaining these data.

***The B-66 and C-133 data were kindly provided by H. Campbell.

****The flying-hour sample for the fire control/gunnery parts is 50,000 hours.

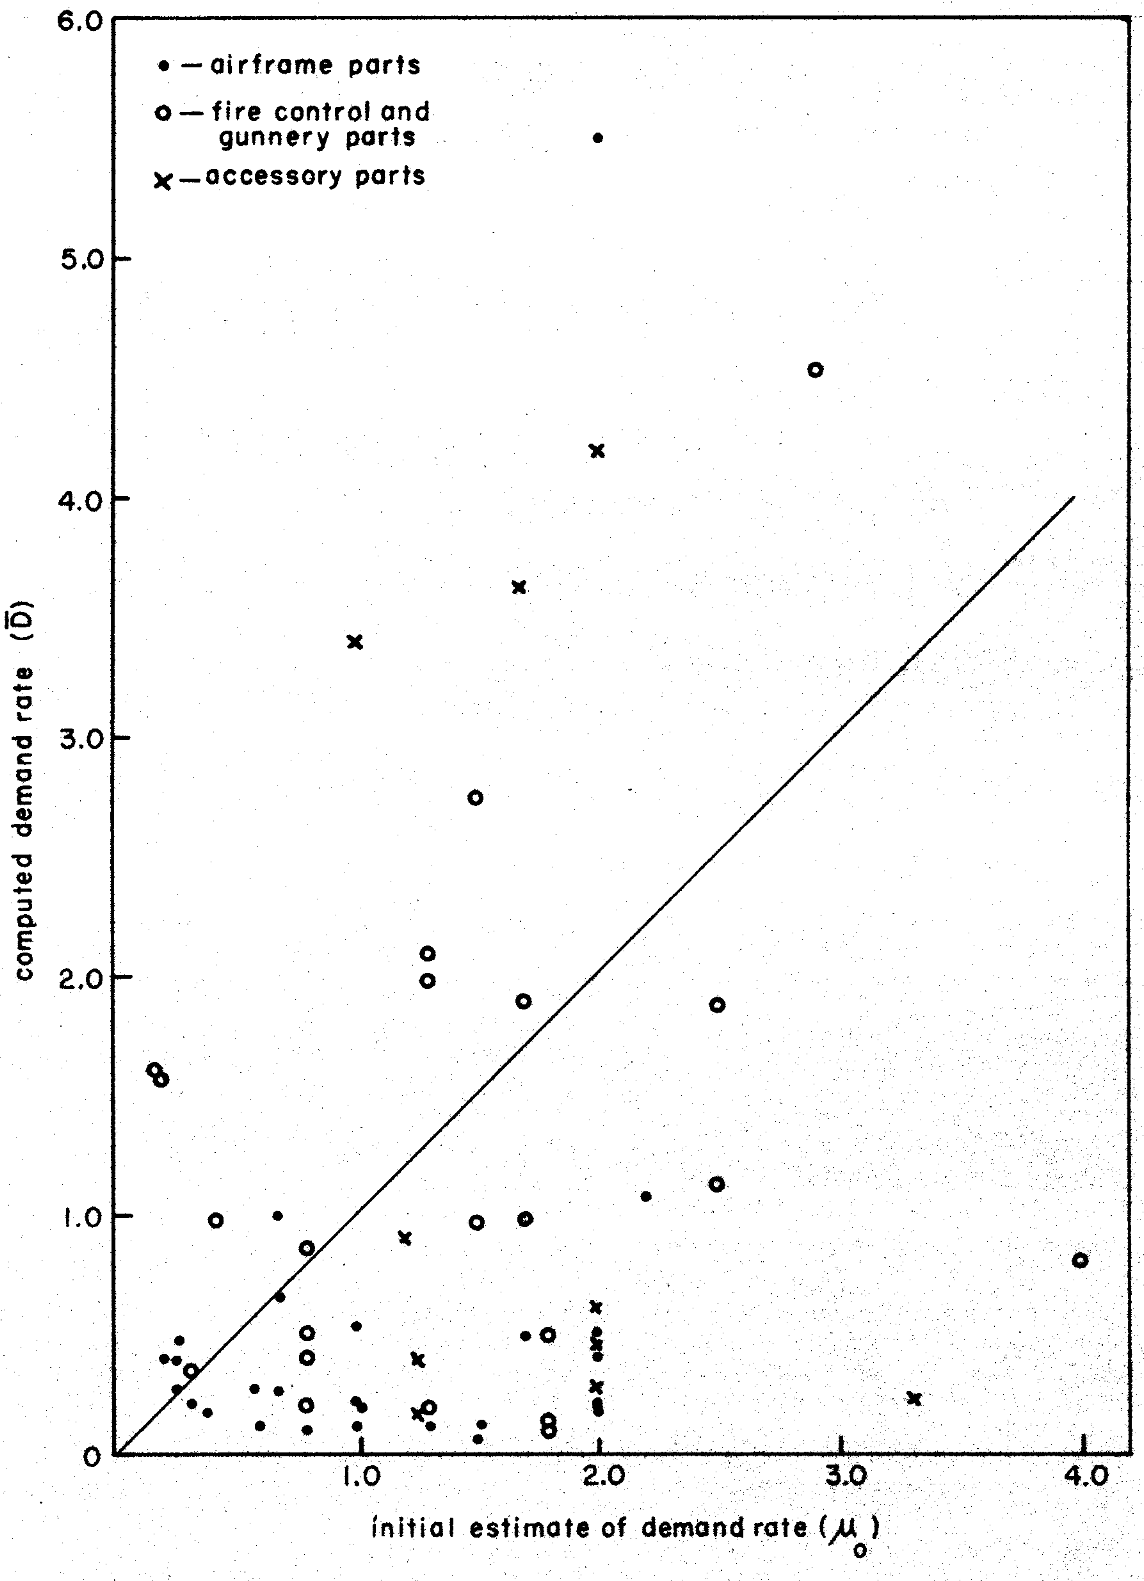

Fig. 5 — Initial estimate versus final computed demand rate for 61 B-52 parts

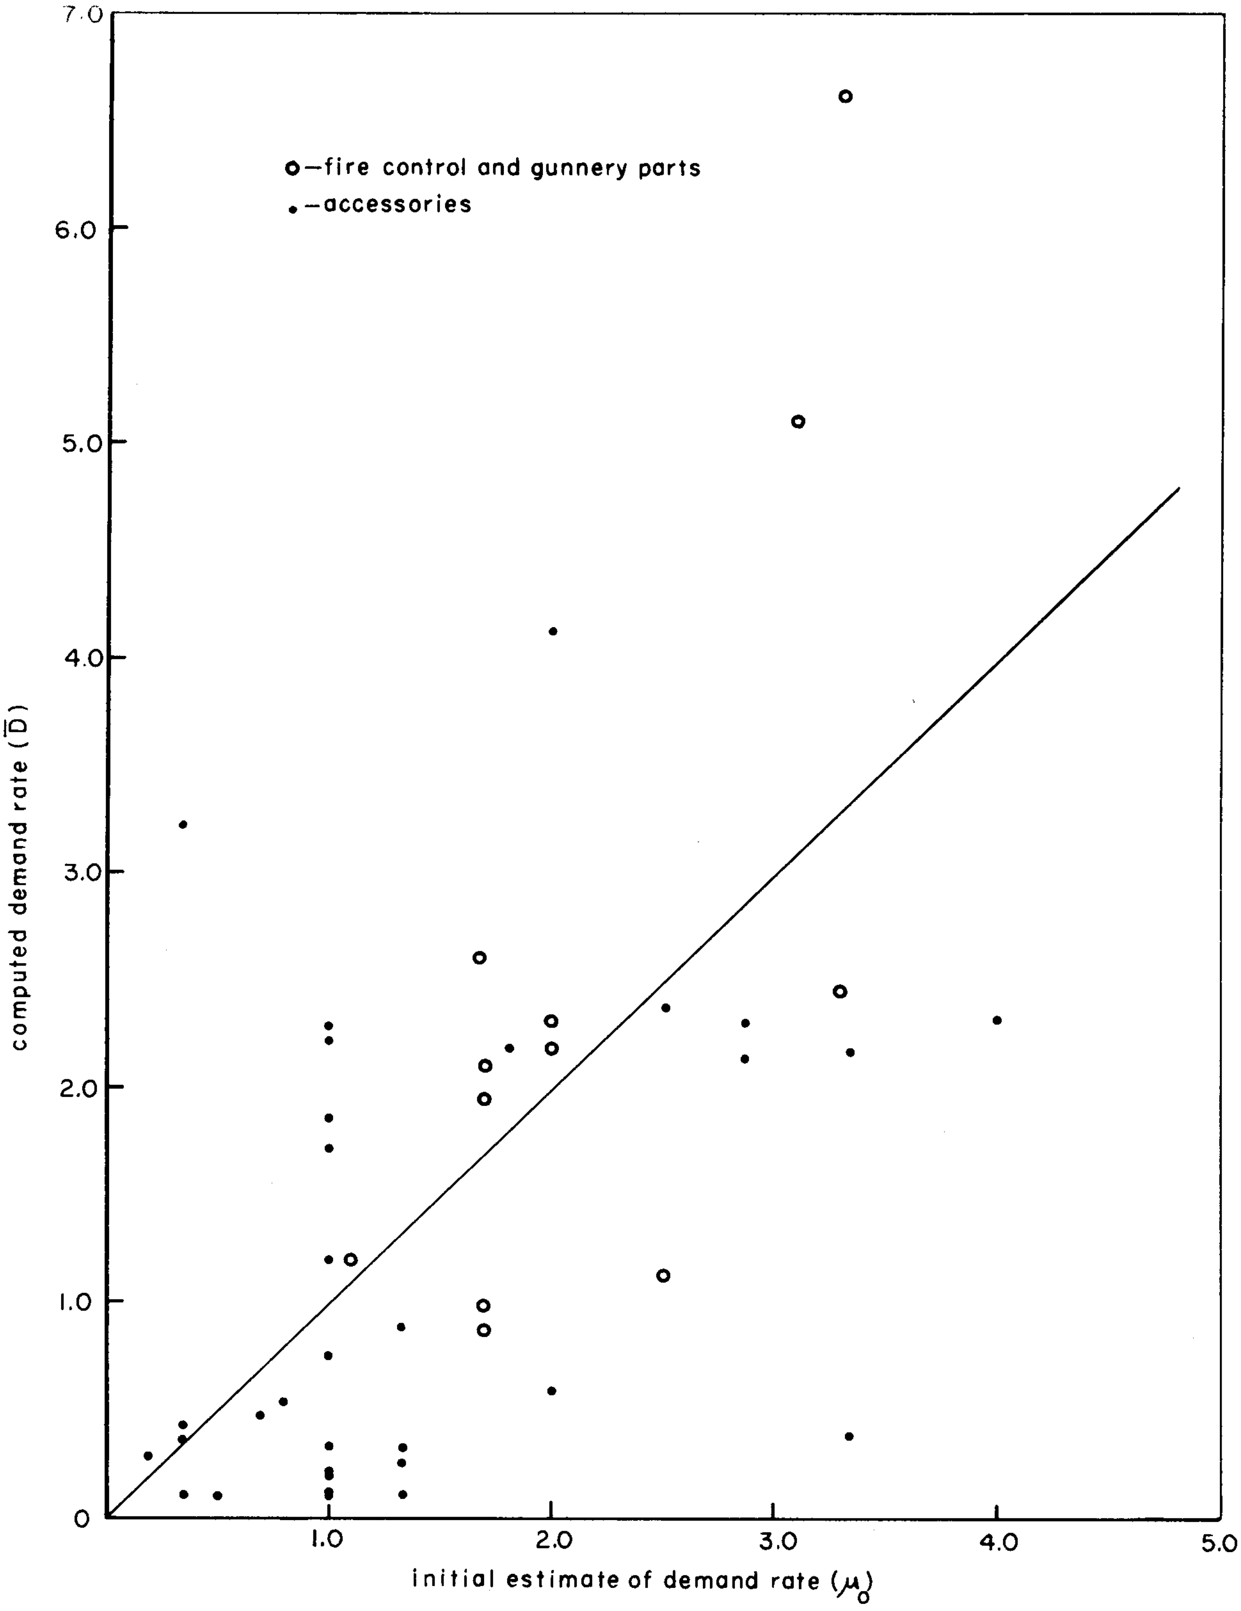

Fig. 6 — Initial estimate versus final computed demand rate for 46 B-66 and C-133 parts

of parts (airframe, fire control/gunnery, and accessories) are plotted separately. The diagonal lines in Figs. 5 and 6 are not intended as a fit of the data, but rather indicate the frequency and extent that the initial estimate proved to be an overestimate (below the diagonal) or an underestimate (above the diagonal).

Figure 6 is a similar diagram for the B-66 and C-133 parts. The computed demand rates for the 19 C-133 parts were based on 18,000 flying-hours; those for the 27 B-66 parts, on a minimum of 20,000 flying-hours.

Figures 5 and 6 show that the initial estimates are not particularly accurate, and that overestimates tend to outnumber underestimates. There is also a strong tendency to overestimate low demand rates and, to a lesser extent, to underestimate high rates. Airframe parts (Fig. 5) tend to be particularly susceptible to overestimates, but this may result from the general tendency to overestimate low-demand-rate parts.

To measure the probable amount of error in μ_0, we need to find the average of the squared differences between μ_0 and the computed rate, $\bar{D}$ (again $\bar{D}$ is computed from a large sample, and is thus a good estimate of the "true" demand rate). For any point in Figs. 5 and 6, the absolute difference between μ_0 and $\bar{D}$ corresponds to the distance between the point and the diagonal, measured either vertically or horizonally. As would be expected, the absolute errors or differences become larger as μ_0 increases. In order to get a meaningful average of the squared differences, we shall need to use the relative, or percentage error, $\frac{\mu_0 - \bar{D}}{\mu_0}$, which is more nearly constant as a function of μ_0. If we eliminate two extreme points (0.16, 1.6) and (0.20, 1.57) in Fig. 5, the average square of the relative differences then becomes

0.69.* After eliminating one point (0.33, 3.23) in Fig. 6, the corresponding value is 0.39. From these limited data, we would recommend that a value of 0.5 be used as the average of the squared relative differences between μ_0 and $\bar{D}$. In practice, we shall need an estimate of the average of the squared absolute differences, σ_0^2, for a given value of μ_0. Using the value recommended above, $\sigma_0^2 = 0.5\,\mu_0^2$.

Further attempts should be made to measure σ_0^2 empirically. In particular, if the estimator indicated high or low confidence in the estimate, we could establish corresponding estimates of σ_0^2. In the meantime, the equation for σ_0^2 given above will serve as a rough approximation of the accuracy of μ_0 for estimating the "true" demand rate.

VALUE OF R

We now return to the problem stated at the beginning of this Section: determining the value of R to be used in obtaining the weights to be attached to μ_0 and $\bar{D}$. Recall from Sec. III that R represents the information-value of μ_0 expressed as an equivalent amount of operational experience.

We have determined empirical measures (s^2 and σ_0^2) of the error distributions of $\bar{D}$ and μ_0. Both s^2 and σ_0^2 are proportional to the expected error

---

*In companion-piece RM-2536, these data were also used to compute an average of the squared relative differences between μ_0 and $\bar{D}$; however, the variable $\dfrac{\mu_0 - \bar{D}}{\bar{D}}$ was used in place of $\dfrac{\mu_0 - \bar{D}}{\mu_0}$. If the data are unbiased, the result will be approximately the same. However, in Fig. 5 the tendency for μ_0 to be much higher than $\bar{D}$ when the latter is small resulted in a grossly different estimate of σ_0^2 when $\bar{D}$ rather than μ_0 was used in the denominator of the above fraction. Since in practice, we shall wish to estimate the average of the squared differences between μ_0 and $\bar{D}$ when μ_0 is given, the variable $\dfrac{\mu_0 - \bar{D}}{\mu_0}$ should be used to measure the relative difference between μ_0 and $\bar{D}$.

when $\bar{D}$ and μ_0 are used as estimates of the "true" demand rate. Recall that the s^2 value is a measure of the $\bar{D}$ error distribution when N=1. If $\sigma_0^2 = s^2$ for a particular part, the informational value of μ_0 is equivalent only to 1000 hours of operating experience. More generally, the ratio s^2/σ_0^2 represents the amount of operational-hours (expressed in 1000's) to which the value of μ_0 is equivalent.* Thus:

$$R = s^2/\sigma_0^2$$

In the previous two Subsections, we have seen that both s^2 and σ_0^2 are estimated as functions of μ_0 . To obtain s^2 we read the value of s in Fig. 3 corresponding to the particular value of $\bar{D}$ (as estimated by μ_0 .) The square of s represents our desired parameter. The value of σ_0^2 , under the recommendation of the preceding Subsection, is 0.5 μ_0^2 . For instance, if μ_0 = 1.0 we obtain an s = 1.6 from Fig. 3; s^2 = 2.56, and:

$$\sigma_0^2 = 0.5(1)^2 = 0.5; \text{ and}$$

$$R = s^2/\sigma_0^2 = 2.56/0.5 = 5.12 \quad .$$

We may similarly obtain the value of R for any μ_0. Figure 7 plots R as a function of μ_0, so that in practice we may determine R immediately from μ_0 without separately computing s^2 and σ_0^2 . Three functions are shown, representing high, medium, and low confidence in the accuracy of the initial estimate. The medium-confidence curve is based on the recommendation of Subsection 3 above; this is the curve we suggest for use in making demand-rate estimates from data currently available in the Air Force. The high-confidence curve might represent a part for which extensive operational experience is

*For a proof of this statement see RM-2536.

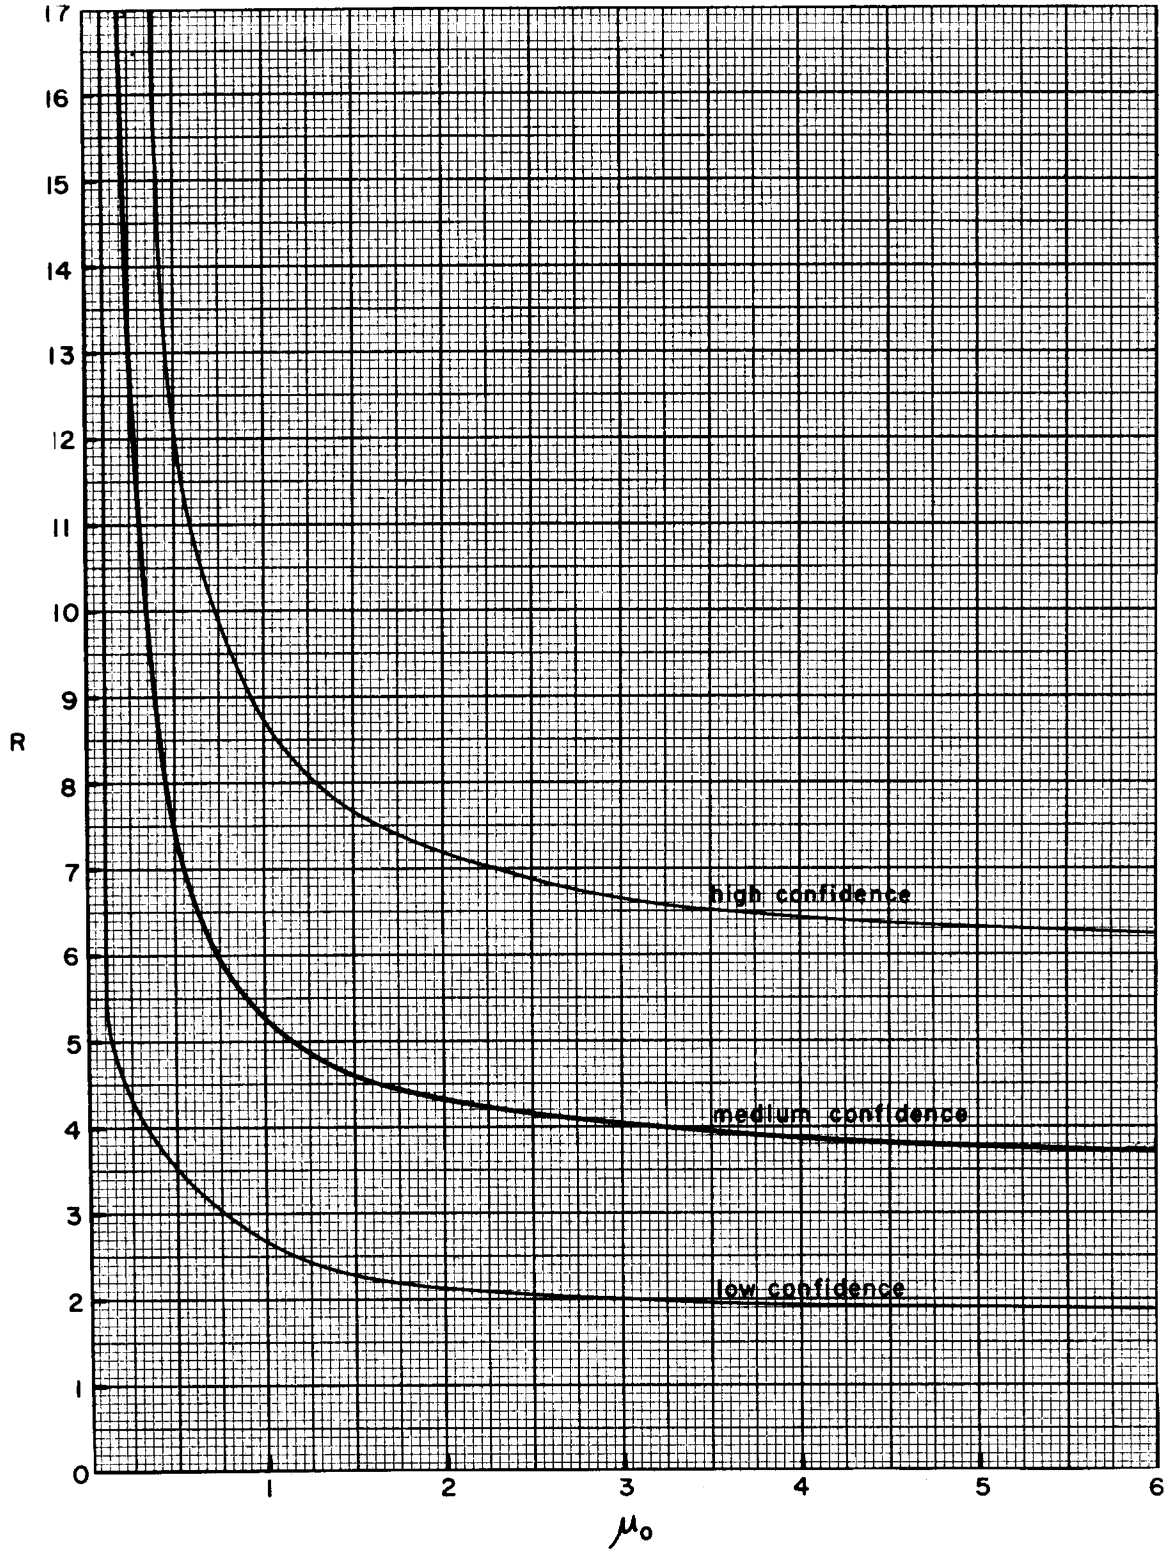

NOTE: Only the medium-confidence curve is based on empirical data; it is the one that should be used with currently available Air Force initial estimates of demand rate.

Fig. 7 — The operational-hour equivalent, R, as a function of the initial estimate, μ_0

available for the identical part used on another weapon. Similarly, the low-confidence curve would represent a part for which there is little or no quantifiable data about the expected demand rate.

V. PROCEDURE

ASSUMPTIONS

Section II discussed the requirement that the past demand-history of a part must reasonably typify future demand experience if it is to be useful for prediction purposes. Accordingly, any differences between future demand rates and those we compute from the sample of operational experience should be random, rather than attributable to part age or changes in the environment in which the part operates. Otherwise, the procedure which follows is not valid for predicting demand rates.* Example 3, given later in this Section, will show how the Bayesian procedure may be used in a limited fashion when the demand rate is not assumed to be independent of part age.

For the Bayesian procedure to be valid, it is obvious that both the initial estimate and the computed rate must refer to the same group of demands. This condition is not met when the initial estimate is intended to include both base and depot reparables, while the observed demand data includes only one of the two.

---

*When the observed demand rate is not representative, because the equipments have undergone scheduled maintenance less than the usual number of times (see the discussion under Case II of Sec. II), we may make an adjustment which will allow us to use the data. We should compute an adjusted number of hours of experience, H', as follows: $H' = H\,[1/k\,(\alpha) + (1 - \alpha)]$,

where $k = \frac{H}{yI}$ and:

- H = total number of hours of operation for all equipment;
- y = number of times equipment underwent scheduled maintenance during H, $y \geq 1$;
- I = interval between scheduled maintenance in hours;
- α = estimate of the proportion of demand for the particular part which normally arises from scheduled maintenance when $k = 1$.

Note that when $\alpha = 1$, the above equation reduces to $H' = H/k$.

STEPS FOR PREDICTING DEMAND RATES

A. Initial Estimate, μ_0

Obtain the estimate of the demand rate expressed in demand per 1000 flying-hours from the AMC 231 form completed at the initial provisioning conference.

B. Size of Operational Sample, N

For aircraft, N should be expressed in thousands of flying-hours. Record N to the nearest tenth; e.g., 2530 flying-hours gives an N of 2.5. For equipment other than aircraft, see Sec. IV, "Error Distribution About $\bar{D}$," p. 19, regarding the unit in which N should be expressed.

C. Computed Demand Rate, $\bar{D}$

Divide the number of observed demands, D, by N:

$$\bar{D} = D/N \ .$$

In determining the number of observed demands, D, any non-recurring demands (see footnote on p. 6) should be eliminated. If some portion of the past operational program is considered non-representative of the future program from the standpoint of demand generation, then both D and N should be revised to include only the representative portion of the data. For example, suppose that 10,000 hours have accumulated with a particular type of aircraft, and that of this total, 3,000 hours resulted from a suitability-test program and 7,000 from routine service operations. The total demand for the 10,000 hours was 12, of which 6 occurred during the test program. If the test program is assumed to be non-representative of service operation, we should revise our data so that D=6 and N=7. The value of $\bar{D}$ would then be 0.86 rather than 1.20.

D. Value of R

Enter Fig. 7 with the value for μ_0 and read the corresponding value of R for medium confidence in the initial estimate.

E. Computation of Weights, w_0 and w_s

$$w_0 = \frac{R}{R + N}\ , \quad w_s = \frac{N}{R + N}\ ; \text{ or,}$$

$$w_s = 1 - w_0\ .$$

F. Predicted Demand μ_1

$$\mu_1 = w_0 \mu_0 + w_s \bar{D}\ .$$

This represents our best prediction of the future demand rate. It is determined jointly by the initial estimate and the computed demand rate, both weighted according to their prediction value.

G. Subsequent Prediction of Demand Rate

After additional operational data become available, the Bayesian procedure can be used to make subsequent demand-rate predictions in the same manner as described above. Under A above, the initial estimate would remain μ_0. Under B, the size of the operational sample would be a new total N. Under C, the new value of $\bar{D}$ would be computed by dividing the total number of observed demands by the total N. Under D, the value of R would be the same. Under E, the weights w_0 and w_s would be found by substituting the new value for N. Under F, the value of μ_1 would be found by substituting the new values for w_0, w_s, and $\bar{D}$.

When N becomes so large that the value of w_0 is no more than 0.10, the

Bayesian procedure should be dropped. $\bar{D}$ alone will then be the predicted demand rate.

CONFIDENCE INTERVAL FOR μ_1

The procedure just described yields our best estimate of the value of the "true" demand rate. This estimate is μ_1. In some cases, we may wish to base procurement and other supply and maintenance decisions on this estimate. When N becomes quite large, the estimate is so close to the "true" rate (granted there is no systematic change in the underlying demand processes) that they may be considered the same. When N is small, however, the "true" demand rate may be appreciably different from μ_1; in this case we may wish to confine ourselves to stating that 90 times out of a 100, let us say, the "true" rate will fall somewhere in the interval $\mu_1 - x$ to $\mu_1 + x$.* We may then wish to base our decisions on a predicted demand rate of $\mu_1 + x$.

When we base our decisions on some demand-rate prediction higher than μ_1, we are buying protection against the chance that μ_1 is actually less than the "true" demand rate. The question of how much protection we should buy is a very complicated problem. Some of the more obvious relevant factors are the cost of the part, the estimated demand rate, manufacturing lead time, and the extent to which flexible maintenance capacity may counteract supply shortages. These problems are beyond the scope of this paper. A formula will be provided here for computing an approximate confidence interval about μ_1, plus two possible alternatives for employing this confidence interval in making

*Note that throughout this discussion we are concerned with a confidence interval about a predicted demand rate or average demand. No probability statement is made concerning the actual number of demands for a given period of time being contained within certain limits. This is a separate problem and concerns the various stockage-level policies.

demand-rate predictions. However, the user is urged to take advantage of any additional information he may have, such as the cost of the part, in deciding whether he should use a predicted demand rate larger than μ_1, and if so, how much larger.

To compute a confidence interval about μ_1, we first enter Fig. 8 with μ_1 (not μ_0), read the corresponding value of s^2,* and then compute the standard error of estimate, σ_1, as follows:

$$\sigma_1 = \sqrt{\frac{s^2}{R + N}}$$

We may then state that approximately 90 per cent of the time the "true" demand rate is less than $\mu_1 + 1.3\ \sigma_1$.**

One decision rule might then be to base procurement and other decisions on a predicted demand rate of $\mu_1 + 1.3\ \sigma_1$. When N is small (less than 5), however, this level of protection may require a substantial investment. If σ_1 is approximately equal to μ_1 for all parts, then it implies that the total number of spare parts purchased (summed over all parts) will be more than twice the number required to support the system. In the above equation we see that σ_1 varies inversely with $\sqrt{N}$; thus as N becomes larger, the cost of the same protection against μ_1 being in error becomes smaller.

This suggests an alternative decision rule to the effect that we should choose as our predicted demand rate the smaller of two quantities: $\mu_1 + 1.3\ \sigma_1$ and $k\ \mu_1$. A reasonable value of k might be 1.5. This rule says we should be

---

*Figure 8 is based on the same data as Fig. 3, but for convenience s^2 instead of s is plotted on the ordinate.

**This statement is only approximate and some liberties have been taken in its derivation, but it is sufficiently accurate for our purposes. If σ_1 is greater than μ_1, no statement should be made concerning the confidence interval.

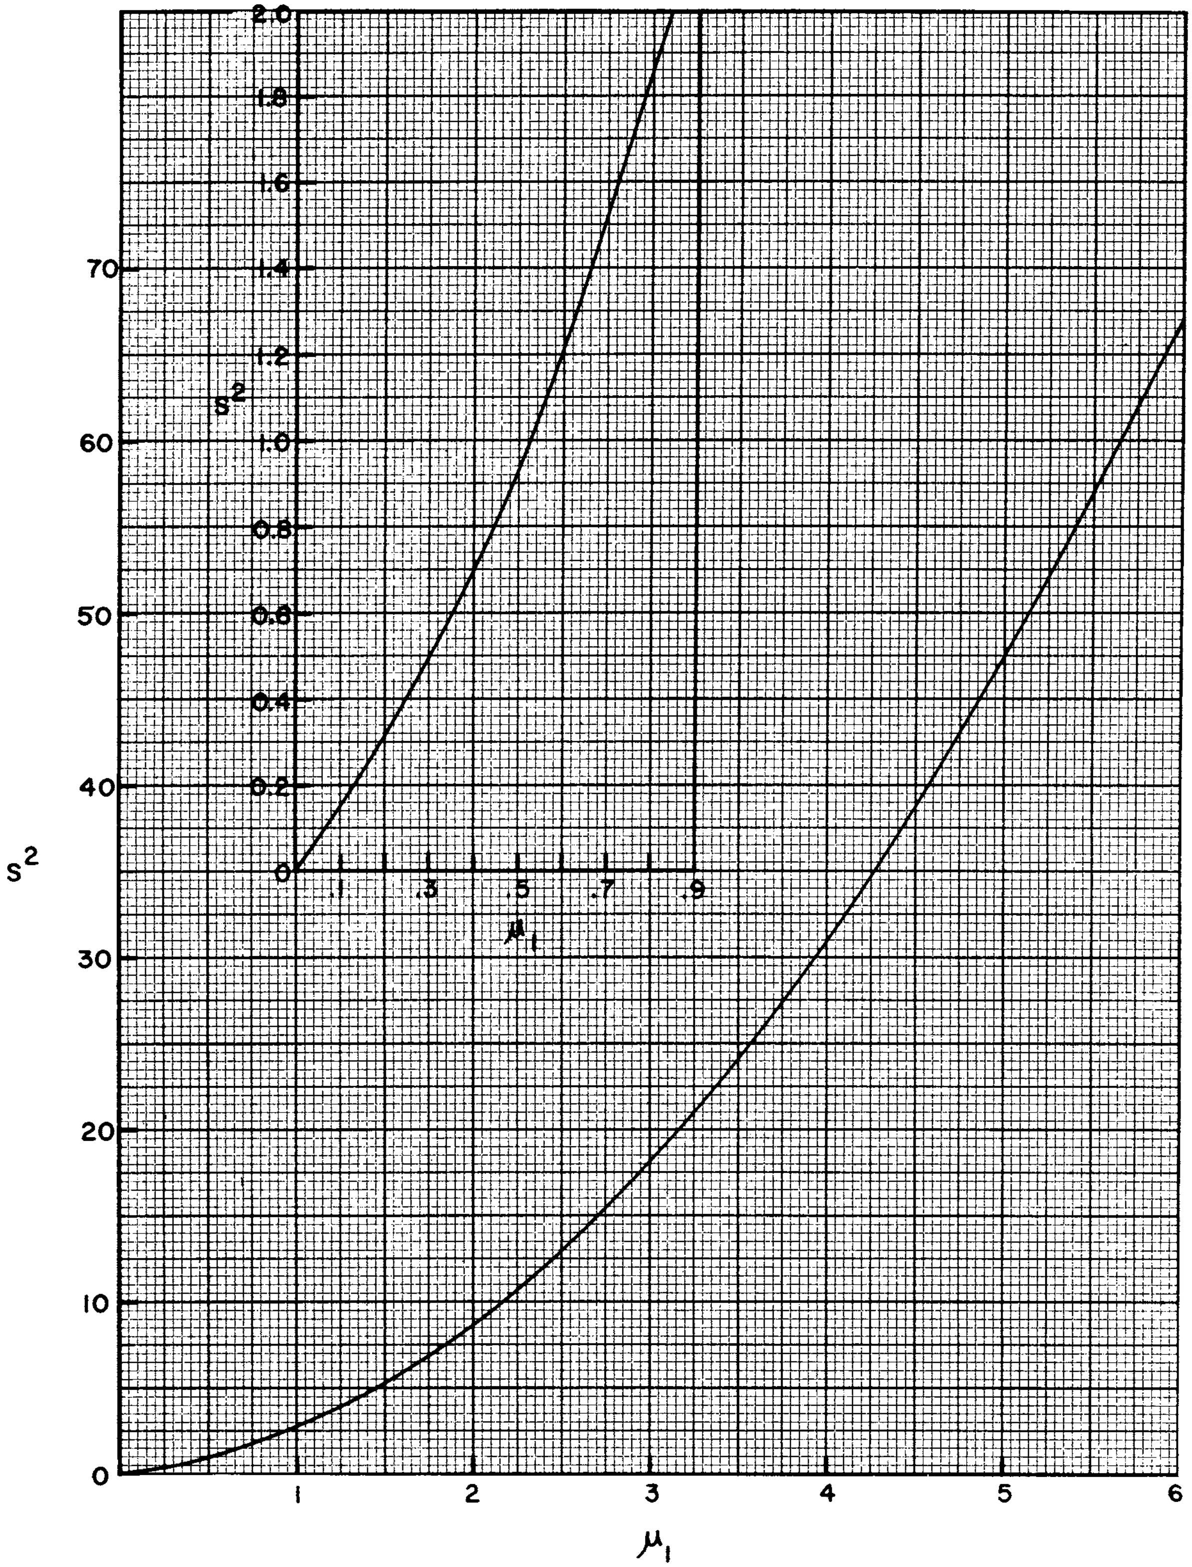

Fig. 8—s^2 as a function of μ_1

[redacted]

90-per-cent certain our prediction is not an underestimate, provided we do not have to pay over a certain amount for the protection.

EXAMPLES

Three examples will illustrate the procedures. In Example 1, the initial estimate of the demand rate for Part X was 1.0 for 1000 flying-hours. After 5000 hours of equipment operation the observed demand was 2. The value of $\bar{D}$ is therefore 0.4. What should be the revised estimate of the demand rate?

Using the medium-confidence curve in Fig. 7, we find an R of 5.2 corresponding to a μ_0 value of 1.0.

The weights w_0 and w_s are then:

$$w_0 = \frac{R}{R + N} \; ; \qquad w_s = \frac{N}{R + N}$$

$$w_0 = \frac{5.21}{5.21 + 5} = 0.51 \qquad w_s = \frac{5}{5.2 + 5} = 0.49$$

The new estimate of the demand rate is:

$$\mu_1 = w_0 \mu_0 + w_s \bar{D}$$

$$= 0.51 \ (1.0) + 0.49 \ (0.4) = 0.71$$

If we wish to employ a confidence interval about μ_1, we first compute σ_1. From Fig. 8 we find that the value of s^2 corresponding to a μ_1 of 0.71 is 1.55; accordingly:

$$\sigma_1 = \sqrt{\frac{s^2}{R + N}} = \sqrt{\frac{1.55}{5.21 + 5}} = 0.39 \ .$$

We may now use μ_1 as our best estimate of the future demand rate for Part X, or we may protect against the chance that μ_1 may be lower than the

"true" rate by employing an upper confidence-interval. We may make the statement that the probability is approximately 0.9 that the "true" demand rate does not exceed $\mu_1 + 1.3\ \sigma_1 = 0.71 + 0.51 = 1.22$.

Suppose that at a later date we have the demand experience for an additional 10,000 flying-hours, and the demand for this period was 3. The value of $\bar{D}$ for an observed demand of 5 and an N of 15 is now 0.33. We simply repeat the procedure outlined above, still using 1.0 as the value of μ_0:

$$w_0 = \frac{5.21}{5.21 + 15} = 0.26 \ ; \quad w_s = \frac{15}{5.21 + 15} = 0.74 \ ;$$

$$\mu_1 = 0.26\ (1.0) + 0.74\ (0.33) = 0.50.$$

To compute σ_1 we enter Fig. 8 with μ_1 and find that $s^2 = 0.90$; thus:

$$\sigma_1 = \sqrt{\frac{0.90}{5.21 + 15}} = 0.21 \ .$$

We may now state that the probability is approximately 0.9 that the "true" demand rate does not exceed $\mu_1 + 1.3\ \sigma_1 = 0.77$. Consequently, our demand-rate estimate should not exceed that value, unless we want even greater protection against an underestimate.

In Example 2, the computed demand rate for Part Y is 2.0 per 1000 flying-hours after 25,000 hours of demand experience. An engineering change is performed on all the Y parts in the inventory, and the contractor predicts that the change will reduce the demand rate for the part by 50 per cent. The customer, however, chooses to estimate the demand rate from operational experience prior to the change, but if subsequent experience substantiates the contractor's prediction, he wishes to revise his estimate. After 10,000

[redacted]

hours of operation subsequent to the change, the demand for Part Y has been 12. In this case, the initial estimate is simply the computed demand rate prior to the change: $\mu_0 = 2.0$. The $\bar{D}$ is $12/10 = 1.2$, and $N = 10$. The value of R for a μ_0 of 2.0 is found in Fig. 7 to be 4.29.*

$$w_0 = \frac{4.29}{4.29 + 10} = 0.30 \quad ; \quad w_s = \frac{10}{4.29 + 10} = 0.70 \; ;$$

$$\mu_1 = 0.30\ (2.0) + 0.70\ (1.2) = 1.44 \text{ ; and}$$

$$\sigma_1 = \sqrt{\frac{4.8}{4.29 + 10}} = 0.58 \; .$$

Our best estimate of the demand rate for Part Y is now $\mu_1 = 1.44$; and $\mu_1 + 1.3\ \sigma_1 = 2.19$. Whether the customer wishes to switch his demand rate estimate to μ_1, or await further operational experience, will depend on the level of protection he wishes to purchase against an underestimate.

Example 3, to be discussed shortly, deals with the case of an age-dependent demand process (see Case V of Sec. II); before proceeding with the example, we should examine this case a little more fully.

An Air Force group has recently advocated the method of an a priori assumption of an age/demand function for high-value parts provisioning.** The provisioning personnel would be asked to select one of several histograms which will express the percentage of the parts they expect to fail during five age intervals. Figure 9 depicts three of these histograms.

---

*We have assumed that the predictive value of μ_0 in this case is the same as if it were obtained in the usual manner.

**Air Force Spares Study Group, Headquarters, Air Materiel Command, Report No. 9, Selective Management of Air Force Materiel in Perspective: 1959, December, 1959.

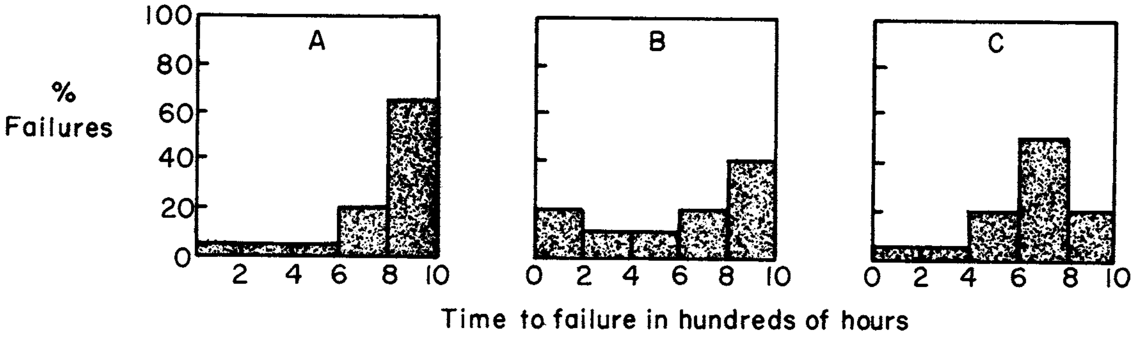

Fig. 9—Failure histograms

These three patterns assume that no part survives more than 1000 hours. The first figure, A, is an expected demand pattern for a part with a mandatory-removal requirement during the 800 to 1000-hour interval. The B pattern suggests a moderately high early demand probability (burn-in), followed by a lower one, and then another increase as the parts age (wearout). The C pattern shows a majority of the parts failing or wearing out around the 600 to 800-hour interval. These patterns are especially useful when definite prior knowledge is available about long-interval scheduled-maintenance policies which will affect the part. This knowledge includes information on mandatory-change requirements; parts not normally demanded except in the overhaul of a higher assembly, such as an engine; and policies which call for returning an assembly (usually electronic) to the depot for overhaul after a given number of repairs or adjustments have been made at base level.

There is some difficulty in using this method when the pattern is based on the hypothesis that a part becomes more susceptible to failure as it ages, rather than on a known scheduled-maintenance policy. In this situation, there is little possibility of revising estimates on the basis of early demand data.

New equipment normally phases into the operational inventory with positive acceleration. A typical pattern for 750 aircraft might be 50 at the beginning of the first year, 200 the second, and 500 the third. Assume that each aircraft averages 30 hours of flying a month, and a particular part is expected to have the age/demand pattern of C in Fig. 9. At the end of the first year we would have the demand experience resulting from 18,000 flying-hours, and at the end of the second year, from 108,000 hours. However, only 6,000 hours of this experience would be applicable to the 600 to 800-hour age interval which, according to the initial hypothesis, should embrace the bulk of the expected failures. Actually, we would have somewhat less than 6,000 hours of experience, since presumably not all of the parts of the first 50 aircraft would have survived to the beginning of the 600 to 800-hour age interval.

The Bayesian weighting procedures may be used here to revise the initial prediction of the demand rate for this interval. To do so, demand must be expressed as a failure rate for those parts which have survived 600 hours of operation,[*] rather than as the percentage of the initial sample expected to fail during that interval. The demand rate within the interval would be considered constant, and, of course, time-referenced failure data would be required to determine how many failures occurred. During phase-in, the age of failing parts can sometimes be approximated if the base where they are used is known, since a base usually receives its full complement of equipment over a relatively short time span.

In Example 3, Part Z is used on an aircraft type that has been operational for three years, and a total of 60,000 hours have been flown. It was

[*]Technically called the "hazard rate."

[redacted]

initially expected to have an age-dependent demand pattern similar to the "C" histogram in Fig. 9. The following table lists the expected demand rates along with the observed demands as a function of the age interval.

Age Interval	Hours Flown	Demand	Demand Rate/1,000 Hrs	
			Expected	Actual
0-200	50,000	7	0.25	0.14
200-400	30,000	3	0.25	0.10
400-600	15,000	4	1.0	0.27
600-800	5,000	0	2.0	0.00
Total	100,000	14	0.45	0.14

The first three intervals have supplied enough experience to get a good approximation of the "true" demand rate. Since the computed rates for these intervals are considerably below those initially predicted, we may wish to revise the expected demand rate for the 600 to 800-hour age interval solely on the basis of the extensive experience in the first three intervals. Suppose we revise it downward to 1.0. We now wish to re-evaluate it on the basis of the sample of 5000 hours obtained in the 600 to 800-hour interval. We assume that the demand rate in that interval is constant, and employ the usual procedures with $\mu_0 = 1.0$; $\bar{D} = 0$; $N = 5$.

$$w_0 = \frac{5.21}{5.21 + 5} = 0.51; \quad w_s = \frac{5}{5.21 + 5} = 0.49 ;$$

$$\mu_1 = 0.51\ (1.0) + 0.49\ (0) = 0.51; \text{ and}$$

$$\sigma_1 = \sqrt{\frac{0.95}{5.21 + 5}} = 0.30 .$$

Our best estimate of the demand rate is:

$$\mu_1 = 0.51; \text{ and } \mu_1 + 1.3\ \sigma_1 = 0.51 + 1.3\ (0.30) = 0.90\ .$$

Our revised estimate of the demand rate during the 600 to 800-hour age interval should therefore be no greater than 0.90.

The use of the Bayesian prediction technique in age-dependent demand processes is not limited to the failure-histogram approach of Fig. 9. Specifically, there is no necessity to assume that all the parts have failed by the time they reach a given age. We may simply predict in advance the demand rate for a given age bracket and then follow the procedure outlined in Example 3 above to revise this estimate with the aid of operational data.

VI. CONCLUSION

This Research Memorandum has described a practical method for predicting spare-parts demand rates, based on the Bayes' Theorem of classical probability theory; it may be readily implemented with a few simple, manual computations. (The detailed procedures and necessary graphs for aircraft parts have been presented in Sec. V.) The method furnishes a quantitative and dynamic approach to demand prediction, by which it is possible to take advantage both of initial estimates and of unfolding operational experience -- even when this experience would be dismissed as too "limited" under current Air Force policy.

The procedure is intended to replace the qualitative approach of assessing a computed demand rate as either "reliable" or "unreliable" depending on the size of the sample. The more quantitative approach described here assigns weights to the initial estimate and to the computed rate which reflect their relative worth for predicting the future demand rate. If there are good reasons for regarding the sample of demand experience as unrepresentative (because of a "burn-in" phenomenon, for example, or part failures that are age-dependent), the computed demand rate should still be disregarded, or at least modified; this rate should not be disregarded solely because the sample size is "limited," however.

The potential user may justifiably complain that he usually has little basis for deciding whether his sample of operational experience is or is not representative of what is to be expected in the future. This question was discussed in Sec. II, which also offered some suggestions for adjusting demand data to make them more representative. In addition, Sec. V presented an example of the use of the Bayesian procedure under conditions of [redacted]

age-dependent demand processes. It was not within the province of this study, however, to try to explain how to identify parts which are subject to these processes.

The utility of the Bayesian technique is not restricted to aircraft parts; it may also be applied to those in missiles or other equipment. The graphs in Figs. 7 and 8 are intended for aircraft, however, and will probably require modification when some operational unit is used other than thousands of flying-hours.

Concerning the confidence in or expected accuracy of the initial estimate of the demand rate, only the curve labeled "medium confidence" in Fig. 7 is based on empirical data. It is this curve that we recommend for application to initial estimates currently available in the Air Force. The high- and low-confidence curves in Fig. 7 are shown to give an idea of how we might treat initial estimates differently, according to how much information they are based on. Should such a classification by confidence rating become available, we should attempt to establish an empirical basis for these two curves as well. Even without a breakdown by confidence levels, additional empirical evaluation of the error distribution around μ_0 would be helpful, since the measurement of σ_0^2 in Sec. IV is based on rather meager data.

The Bayesian procedures described here should thus be especially helpful when fairly large commitments must be made for spare parts with the help of only a small sample of demand experience, as they often must with rapidly obsolescing military systems.