

**NAVAL AEROSPACE MEDICAL RESEARCH LABORATORY
51 HOVEY ROAD, PENSACOLA, FL 32508-1046**

NAMRL Special Report 00-2

**A COST-BENEFIT ANALYSIS OF THE LANDING CRAFT AIR CUSHION
(LCAC) SELECTION SYSTEM**

D. J. Blower

ABSTRACT

The purpose of this report is to document a cost-benefit analysis of the Landing Craft Air Cushion Vehicle (LCAC) Selection System for craftmasters and engineers. A cost-benefit analysis for this selection system had not been conducted before, and it seemed worthwhile to see if there was a cost justification for the continued use of this system. The analysis in this paper indicates an annual net savings somewhere in the range of no savings to \$350,000. The best guess is an annual net savings of about \$160,000. About 70% of the distribution is centered on the range of \$60,000 to \$260,000 net savings per year. Because the bulk of the distribution covers an expected cost benefit to the LCAC training commands, we recommend the continued usage of the LCAC Selection System to prefilter candidates for training as craftmasters and engineers. Monitoring of the data and updates to the cost structure should be carried out periodically to determine if these savings can be expected to continue into the future.

Acknowledgments

I would like to acknowledge the congenial work environment provided by LT Henry P. Williams, Head of the Selection Division at the Naval Aerospace Medical Research Laboratory. In addition, thanks are due to LT Rick Arnold of the Naval Operational Medicine Institute for providing the latest data on which this report is based. As always the support provided by Dr. Ken Ford of the Institute for Human and Machine Cognition at the University of West Florida is gratefully acknowledged.

INTRODUCTION

The ultimate justification for selection systems in the military is to help reduce training costs. By prefiltering candidates before they enter training, selection systems lower the actual attrition rate that would have been evident had the selection system not been in place. If this difference in attrition rates with and without a selection system is large enough, then there are quantifiable cost savings to the training budget.

The purpose of this report is to document a cost-benefit analysis of the Landing Craft Air Cushion Vehicle (LCAC) Selection System for craftmasters and engineers. A cost-benefit analysis for this selection system had not been conducted before, and it seemed worthwhile to see if there was a cost justification for the continued use of this system.

Many elements of the cost-benefit analysis are subject to uncertainty. In the present circumstances, we are uncertain about 1) the baseline attrition rate for LCAC craftmasters and engineers if a selection test battery were not operational, 2) the reduced attrition rate after the selection test battery has been put into place, 3) the actual costs for each training attrition, 4) the costs associated with replacing candidates who are rejected by the test battery, and 5) the number of students that will be trained in any given year.

Despite these uncertainties, empirical data gathered from the use of the LCAC Selection System over the past few years and some reasonable estimates of the costs outlined above can be leveraged to construct a distribution of savings. The optimal way of handling these uncertainties is through probability theory. The Bayesian approach to data analysis uses probability theory to make the best inference conditioned on known information. In this paper, the Bayesian predictive distribution is used to help answer the question of whether the LCAC selection system reduces the attrition rate during training.

The analysis in this paper indicates an annual net savings somewhere in the range of no savings to \$350,000. The best guess is an annual net savings of about \$160,000. About 70% of the distribution is centered in the range of \$60,000 to \$260,000 net savings per year. Because the bulk of the distribution covers an expected cost benefit to the LCAC training commands, we recommend the continued usage of the LCAC Selection System to prefilter candidates for training as craftmasters and engineers. Monitoring of the data and updates to the cost structure should be carried out periodically to determine if these savings can be expected to continue into the future.

COSTS AND PROBABILITIES

According to the LCAC training community, it costs \$160,000 to train a craftmaster or engineer for the initial 17 week training period. Approximately 30 students are trained in any given year. Therefore, we may take the annual training budget to be \$4,800,000. When the LCAC community first approached NAMRL in 1987 to help reduce training attritions, the probability of an attrition was as high as 40%. The empirical data over the past 13 years indicate that the probability of an attrition has ameliorated from that initial high level to somewhere in the range of 20 to 30%. One reasonable assignment of the probability for an attrition if there were no selection test battery is in the middle of this range at 25%. The cost due to attrition is thus

$$\$4,800,000 \times .25 = \$1,200,000.$$

With the test battery operating to screen out potential attritions, one reasonable estimate as to the revised probability of attrition is about 18%. The justification for such a number is provided later in the report. The cost due to this revised attrition is thus

$$\$4,800,000 \times .18 = \$864,000.$$

Therefore, the savings based on this difference in attrition rates is estimated at

$$\begin{array}{r} \$1,200,000 \\ - \quad \$864,000 \\ \hline \$336,000 \end{array}$$

This savings is attributable to the effectiveness of the selection test battery in lowering attrition rates. We shall label this difference as Δ , (delta). So, in this first example, we can find the savings just by computing a delta of

$$\Delta = .25 - .18$$

$$= .07$$

$$.07 \times \$160,000 \times 30 = \$336,000.$$

There are costs attached to the use of the LCAC Selection System, so these must be subtracted from the savings just calculated to arrive at a net savings due to the system. For the purpose of this report, we list three such administrative costs. We assign \$50,000 for the travel and per diem costs to transport prospective LCAC trainees to the Naval Operational Medicine Institute (NOMI) in Pensacola, Florida. This figure can be ascertained fairly accurately since NOMI tests about 60 candidates per year, and the average travel and per diem costs are about \$850 per candidate.

The second cost concerns routine administration and upkeep of the LCAC test battery. NOMI must allocate personnel to run the test battery, maintain data bases, and oversee its administration. The system must be calibrated, checked, and undergo periodic software and hardware upgrades. For all of this, we assign an arbitrary cost of \$40,000 per year.

Finally, the LCAC Selection System will reject some percentage of the candidates sent to Pensacola. They will be rejected because the Selection System predicts them as failures during training. Currently, our best guess is that the system will reject about 38% of the candidates tested.¹ This is the most difficult cost to assess. How expensive is it to replace those candidates rejected by the test battery? If there is a large pool of qualified applicants, then this cost must be less than if there is difficulty in recruiting volunteers. For the sake of conducting these numerical exercises, \$35,000 is assigned for this cost.

These costs are simply my best guess so that I could commence with the numerical examples. I welcome the experts in the LCAC training community to critique these costs and provide more realistic numbers should they exist. However, the techniques for assessing the merit of the LCAC Selection System as outlined in this report remain the same. Any better cost estimates can be substituted into the framework provided here, and new analyses can easily be run to judge their impact. With these estimates in place, the net savings ascribable to the LCAC Selection System for this example can be calculated to be \$211,000.

	\$336,000
-	\$ 50,000
-	\$ 40,000
-	\$ 35,000
	\$211,000

BRACKETING THE EXPECTED NET SAVINGS

In the Introduction, it was mentioned that the uncertainty about the net savings could be bracketed between no savings at the low end and close to \$400,000 at the high end. The no savings at the low end results from a set of very pessimistic assumptions, while the savings at the high end results from a set of very optimistic assumptions. We will eventually argue that the truth lies somewhere between these extreme sets of assumptions. The set of pessimistic assumptions is examined first.

Contrary to the initial example given above, suppose that the true rate of attrition without the candidates first going through the test battery is not 25%, but rather a lower value of 20%. And further, under this set of

¹ This figure is subject to change because the threshold score needed for a predicted pass was lowered in July 1998.

pessimistic assumptions, suppose that the LCAC test battery provides no extra information about a candidate's chance for success during training. With no extra information from the candidate's score on the test battery, the failure rate for the LCAC selected candidates remains at 20%. In this case, $\Delta = 0$, and there is no savings at all due to the difference in attrition rates. The "net savings" is actually a loss of $-\$125,000$ due to the costs associated with operating the test battery. Thus, it costs more to have a selection test battery than if candidates skipped the entire process and went directly into training.

One the other hand, one could indulge in a very optimistic set of assumptions to arrive at a markedly different conclusion. Under this set of assumptions, the true rate of attrition rises to 30% without the selection test battery. The test battery, in addition, is actually more powerful in weeding out unsuccessful candidates than the limited sample size has led us to believe. Suppose that the rate of attrition when candidates are first screened by the test battery is only 12%. In this fortunate case, $\Delta = .30 - .12 = .18$. The net savings realized is

$$\begin{aligned}\text{Savings} &= .18 \times \$160,000 \times 30 \\ &= \$864,000 \\ \text{Net Savings} &= \$864,000 - \$125,000 \\ &= \$739,000.\end{aligned}$$

Of course, neither of these extreme set of assumptions is likely to be the truth. The truth lies somewhere in the middle. That is why under a more reasonable set of assumptions, the extreme values do not lie between $-\$125,000$ and $\$739,000$, but rather lie between the more restricted numbers given in the introduction. We now turn to examine the data on which to base this reasonable set of assumptions.

ORGANIZATION OF THE FREQUENCY DATA

The estimates for the baseline attrition rate, that is, the attrition rate without a selection system, and the adjusted attrition rate after the implementation of a selection system, are based on empirical data. The LCAC selection system has been operational at NOMI for about 8 years, beginning in October 1992. In this report, frequency data are examined from that initial start date to the present. In addition, there are data from the R&D phase prior to October 1992 when validation testing and initial operational usage took place at NAMRL.

These frequency counts are best presented in a 2×2 table as sketched in Fig. 1. The two columns of the table represent the predicted passes and the predicted failures, and the two rows represent the actual passes and the actual failures. The numbers in the Predicted Pass column are the number of candidates who achieved scores above the composite score of $+14$, and the number in the Predicted Fail column are the number of candidates who achieved scores below that composite score. See Blower [3] for a description of the composite scores and the threshold score.

There are four cells in the table that indicate the joint occurrence of one of the rows and columns. The breakdown of the predicted pass-actual pass cell and the predicted pass-actual fail cell is known from the training data. These are candidates who scored above the threshold and who therefore entered training. However, the breakdown of the predicted fail-actual pass cell and the predicted fail-actual fail cell is unknown because these cells represent the candidates rejected by the system. They never entered training and therefore we don't know how they would have fared in training.

Nevertheless, some of the frequency counts in the data base can be placed in these last two cells. The subjects who participated in the R&D phase at NAMRL all entered training whatever their score on the test battery. This entry into training despite the score on the test battery occurred during the validation stage of the selection system. Also, the composite score and the threshold score in the early days of the selection system were based on different weightings and different predictor variables. In 1995, the threshold score was set at a value of $+14$ and remained

	Predicted Pass	Predicted Fail	
Actual Pass	Cell 1 Score above threshold and pass training*	Cell 3 Score below threshold and do not enter training*	Marginal Sums are recorded here
Actual Fail	Cell 2 Score above threshold and fail training*	Cell 4 Score below threshold and do not enter training*	
Marginal sums are recorded here			
Composite Score ≥ +.14		Composite Score < +.14	
* See text for exceptions			

Figure 1: A 2×2 table to organize the empirical frequency data.

there until July 1998. At that time, the threshold score was moved down to $-.34$ and it has remained there until the present.

The $+.14$ threshold score will be used as an arbitrary dividing line to place the data into either the predicted pass or predicted fail column². Therefore, because of the changing way of computing composite scores and threshold scores, there were a few early candidates who scored above $+.14$ but were rejected by the system as predicted fails. There are more candidates who scored below $+.14$, but because of the different conditions in effect at that time, were admitted into training. These candidates fall into the predicted fail column (given our criterion of separation at $+.14$), but since they did, in fact, enter training they can be placed into one of the two cells.

Thus, some subjects in the data base can be unequivocally allocated to one of the four cells while the true status of other subjects remains unknown. These are all the subjects who scored below $+.14$ when that threshold was in effect, and those subjects who scored above $+.14$ when a different threshold was in effect and did not enter training. In addition to these subjects, the data base contains subjects who have taken the test battery and were predicted passes, but who have not yet started training.

We would like to make a reasonable allocation of these subjects whose true status is unknown to one of the four cells of the table. Using all of this information, we can make inferences about the attrition rate with and without the selection system in place. This difference is needed so that Δ can be used to calculate the cost savings.

In one case, we can make an extreme allocation where we place *all* of the predicted fails into the predicted fail-actual fail cell. We also place *all* of the predicted passes who did not enter training into the predicted pass-actual pass cell. This extreme allocation favors the LCAC system to the maximum extent possible. It is what we have labeled as the set of extremely optimistic assumptions above.

In the other extreme allocation, *all* of the predicted fails can be placed into the predicted fail-actual pass cell, and *all* of the predicted passes placed into the predicted pass-actual fail cell. This form of an extreme allocation discredits the LCAC selection system to the maximum extent possible. It is what we have labeled as the set of extremely pessimistic assumptions in the discussion above.

²Any other composite score could have been used as the arbitrary dividing line to separate predicted passes and predicted fails. In fact, we could choose a number of these different threshold scores to trace out an Receiver Operating Characteristic (ROC) curve from Signal Detection Theory.

It is much more likely that there is some "reasonable" split of these subjects into the four cells. We use the Bayesian approach to find such a reasonable split. Specifically, the Bayesian predictive distribution will help to ascertain the probability of various splits among the four cells given the known data.

THE OBSERVED DATA

Now let's look at some of the empirical data. Combining the data from the NAMRL R&D phase and the NOMI operational phase results in the frequencies given in Fig. 2. The first number given in each of the four cells

	Predicted Pass		Predicted Fail	
Actual Pass	NAMRL NOMI 50 + 166 216 Cell 1	NAMRL NOMI 17 + 27 44 Cell 3	260	
Actual Fail	NAMRL NOMI 8 + 35 43 Cell 2	NAMRL NOMI 6 + 13 19 Cell 4	62	
	259	63	322	

Figure 2: A breakdown of frequency counts into a 2×2 table. These counts are known to be correctly placed into one of the four cells.

is the NAMRL data, and the second number is the NOMI data. These numbers are all correctly placed into one of the four cells. The numbers are correctly placed under the predicted pass column because all of these subjects did enter training, and we know their training outcome. The numbers under the predicted fail column are also correctly placed because, although their current threshold scores are below $+14$, at the time they took the test battery a different algorithm was in effect and they were predicted to be passes. They also entered training and we know their outcome as well. This column reflects, as well, those subjects tested at NAMRL during the validation stage who entered training no matter what their composite score.

Also in the data base are subjects whose training status is unknown. There are presently 196 subjects in this category. Forty of these subjects were predicted passes given the dividing line threshold score of $+14$. They can be subdivided into two classes: those awaiting training whose training status will eventually become known and those who never entered training because at the time a different algorithm had them as predicted fails. A total of 156 subjects falls below the threshold score of $+14$ and are predicted fails. These subjects can also be subdivided into two classes. The bulk of these subjects were rejected by the test battery when the algorithm in effect at the time used a threshold score of $+14$. Consequently, they never entered training and we will never know what their training outcome would have been. The threshold score was changed from $+14$ to $-.34$ in July 1998 as mentioned above. Therefore, a small number of subjects will have threshold scores below $+14$ but above $-.35$ so they will enter training, and eventually their outcome will be known as well.

EXTREME ALLOCATION STRATEGIES

To examine the extreme allocation strategies given these data, consider first the extreme allocation strategy most disfavorable to the LCAC selection system. That is, allocate the 40 predicted passes whose status is unknown to the predicted pass-actual fail cell. Then allocate the 156 predicted fails whose status is unknown to the predicted fail-actual pass cell. See Fig. 3 where the first number in each cell is taken from Fig. 2 and the second

number in parentheses is dictated by the allocation strategy. The new marginal totals are also presented. Use these

	Predicted Pass	Predicted Fail	
Actual Pass	216 + (0) 216	44 + (156) 200	416
Actual Fail	43 + (40) 83	19 + (0) 19	102
	299	219	518

Figure 3: An extreme allocation of the 196 subjects whose status is unknown. This allocation is the one most disfavorable to the LCAC selection system.

marginal totals to estimate the attrition rate with and without the selection system. Not surprisingly, under this extreme allocation that biases to the maximum extent possible against the selection system, the difference in attrition rates favors having no selection system. The Δ is about negative 8%.

$$\begin{aligned}
 P(\text{Attrite without selection system}) &= \frac{102}{518} \\
 &= 19.69\% \\
 P(\text{Attrite with selection system}) &= \frac{83}{299} \\
 &= 27.76\% \\
 \Delta &= 19.69\% - 27.76\% \\
 &= -8.07\%.
 \end{aligned}$$

On the other hand, examine the extreme allocation strategy that favors the LCAC selection system to the maximum extent possible. In this case, we just reverse the placement of the 196 subjects whose status is unknown. That is, the 40 subjects who were previously placed into the predicted pass-actual fail cell are now placed into the predicted pass-actual pass cell, and the 156 subjects who were previously placed into the predicted fail-actual pass cell are now placed into the predicted fail-actual fail cell. See Fig. 4 for this new rearrangement. The marginal totals on the right-hand side are affected by this change, but the marginal totals along the bottom are not. We use the new marginal totals to once again estimate the attrition rate with and without the selection system. Not surprisingly, under this extreme allocation that favors the selection system to the maximum extent possible, the difference in attrition rates strongly favors the selection system. The Δ is almost 28%. The savings in this case would be enormous.

$$P(\text{Attrite without selection system}) = \frac{218}{518}$$

	Predicted Pass	Predicted Fail	
Actual Pass	216 + (40) 256	44 + (0) 44	300
Actual Fail	43 + (0) 43	19 + (156) 175	218
	299	219	518

Figure 4: An extreme allocation of the 196 subjects whose status is unknown. This allocation is the one most favorable to the LCAC selection system.

$$\begin{aligned}
 &= 42.08\% \\
 P(\text{Attrite with selection system}) &= \frac{43}{299} \\
 &= 14.38\% \\
 \Delta &= 42.08\% - 14.38\% \\
 &= 27.70\%.
 \end{aligned}$$

Now, no one believes in either of these extreme allocation strategies. What kind of technique can be used to accomplish a more reasonable allocation of these 196 subjects?

THE BAYESIAN PREDICTIVE DISTRIBUTION

There is obviously some uncertainty attached to how we should allocate the 196 subjects with an unknown training status to a known training status. They could be split up in any number of ways. Two ways, albeit seemingly extreme, were just discussed on how to accomplish that split. This was done to bracket all the ways the split could be achieved by the most favorable and the most unfavorable to the cost-benefit analysis of the LCAC selection system.

Intuition would tell us that, not knowing anything else that should influence the allocation, we should follow the ratio of the subjects whose training status is known. The Bayesian predictive distribution does what our intuition tells us should be done, but in a precisely quantifiable manner. The derivation of the Bayesian predictive distribution will not be repeated here. The technical details of the derivation and application to problems can be reviewed in Blower [1,2,4].

Technically, the predictive distribution used here to solve this allocation problem is called the beta-binomial distribution. It shall be a guide to making reasonable allocations of subjects which we sought as an alternative to the extreme strategies. Before we arrive at the technical definition of the predictive distribution, let us state in words what we are doing.

If we know that a penny is fair, then we have no problem determining a "reasonable split" between Heads and

Tails. In 100 tosses of the penny, 48 Heads and 52 Tails would be considered a reasonable split, but 98 Heads and 2 Tails would not if the penny were actually fair. If we didn't know the penny was fair, but had tossed it a number of times in the past and recorded the number of Heads and Tails, we could use this empirical data to predict future outcomes. If we had gotten six Heads in ten previous tosses we would intuitively feel that the probability for Heads could very well lie between say .3 and .7, but not between .05 and .15. We would average over the various probabilities for Heads given this kind of support by the past empirical data in assessing the chances for obtaining a split of 20 Heads and 30 Tails in 50 future tosses.

Let $L(z|\theta)$ stand for the binomial likelihood of obtaining z "successes" in N trials. In our problem, z stands for the number of subjects to allocate to cell 1 (predict pass-actual pass), and $N - z$ stands for the number of subjects to allocate to cell 2 (predict pass-actual fail). Now we introduce a parameter called θ that influences the chance of each individual going into cell 1. Then clearly, $1 - \theta$ is the parameter that influences the chance of each individual going into cell 2. The binomial formula is

$$L(z|\theta) = \binom{N}{z} \theta^z (1 - \theta)^{N-z}. \quad (1)$$

We have two allocation problems. The first is to allocate the 40 subjects who are predicted passes into the first column of the 2×2 table, and the second is to allocate the 156 subjects who are predicted failures into the second column. We address the first allocation problem. N in this case is 40 and z could be anywhere from 0 to 40. The binomial formula gives us another reason to consider the extreme allocation of $z = 0$ or $z = 40$ to be highly unlikely.

Let's re-examine the situation that was maximally unfavorable to the LCAC selection system. In this case, all 40 subjects were allocated to cell 2, so $N - z = 40$ and $z = 0$. For the time being, suppose we begin in a state of initial ignorance about the parameter θ , and since there are only two cells where the subjects could be allocated, we assign $\theta = (1 - \theta) = .5$. Using Equation (1), the likelihood for this allocation of 40 subjects to cell 2 is,

$$\begin{aligned} L(z = 0|\theta = .5) &= \binom{40}{0} .5^0 .5^{40} \\ \binom{40}{0} &= \frac{40!}{0! 40!} \\ &= 1 \\ .5^0 &= 1 \\ .5^{40} &= 9.09 \times 10^{-13} \\ L(z = 0|\theta = .5) &= 1 \times 1 \times 9.09 \times 10^{-13} \\ &= 9.09 \times 10^{-13}. \end{aligned}$$

Similarly, the same low likelihood is obtained in the scenario most favorable to the LCAC Selection System. Now $z = 40$ and $N - z = 0$, but you can see from the symmetry of Equation (1) that this doesn't make any difference. What does make a difference in the likelihood is a more reasonable split. In the current example, the maximum likelihood is obtained by an even split between cell 1 and cell 2. Now $z = 20$ and $N - z = 20$, and the binomial likelihood is

$$L(z = 20|\theta = .5) = \binom{40}{20} .5^{20} .5^{20}$$

$$\begin{aligned}
\binom{40}{20} &= \frac{40!}{20! 20!} \\
&= 1.38 \times 10^{11} \\
.5^{20} &= 9.54 \times 10^{-7} \\
.5^{20} \times .5^{20} &= 9.09 \times 10^{-13} \\
L(z = 20|\theta = .5) &= (1.38 \times 10^{11}) \times (9.09 \times 10^{-13}) \\
&= .1254.
\end{aligned}$$

Here we see that an even split has a much higher likelihood than an extreme split. You are much more likely to obtain 20 Heads and 20 Tails in 40 tosses of a fair coin than no Heads and 40 Tails. For a parameter setting of $\theta = .5$, this is due entirely to the $\binom{N}{z}$ term. An even split of 40 subjects can be accomplished in a vastly greater number of ways as compared to the one way for an extreme split. For a more detailed explanation of this combinatorial argument see the appendix.

In the beginning of this section, we said that the predictive distribution was an averaged likelihood. An average, by definition, is an integral of the object being averaged with respect to a continuous probability distribution. The object being averaged is the likelihood and the probability distribution is for θ . The parameter θ can only take on values between 0 and 1, so the average likelihood is

$$\text{Average Likelihood} = \int_0^1 L(z|\theta) p(\theta) d\theta. \quad (2)$$

The Bayesian twist to this expression is that the probability of θ can be refined by taking into account the known frequencies of falling into these two cells. From Fig. 2 the ratio for cell 1 is 216/259 and the ratio for cell 2 is 43/259 and we might take this as a good guess to help us allocate the subjects with unknown training status. This is just like the penny that we didn't know was fair, but had tossed a number of times and recorded the outcomes. The Bayesian formalism accomplishes this by constructing a posterior distribution for θ based on these known training outcomes under the predicted pass column. Therefore, Equation (2) is amended by inserting the posterior distribution for θ as conditioned on, say, y known successes from n trials. In this case, $y = 216$ and $n = 259$.

The average likelihood of Equation (2) is now called the predictive distribution and written as

$$\begin{aligned}
\text{Average Likelihood} &= \int_0^1 L(z|\theta, N) p(\theta|y, n) d\theta \\
&= P(z|n, y, N) \\
&\equiv \text{Predictive Distribution}
\end{aligned} \quad (3)$$

Using a computer program we can calculate Equation (3), the predictive distribution, for all values of z from $z = 0$ to $z = 40$. Figure 5 shows a graph of this predictive distribution. Since the observed data of known frequencies is skewed to higher numbers in cell 1, the predictive distribution is also concentrated at higher values of z . The mode, the most probable value of z , occurs at $z = 34$ with

$$P(z = 34|n = 259, y = 216, N = 40) = .1560.$$

Neighboring values start to taper off from this maximum value, so that by the time z reaches 27 going in the downward direction from the mode,

$$P(z = 27|n = 259, y = 216, N = 40) = .0102$$




Figure 5: The Bayesian predictive distribution for the allocation of 40 subjects who were predicted passes (but who did not, or have not yet entered training) into the predicted pass-actual pass cell.

and when z reaches 38 going in the upwards direction from the mode

$$P(z = 38|n = 259, y = 216, N = 40) = .0253.$$

The bulk of the probability distribution resides between these two values with

$$P(27 \leq z \leq 38|n = 259, y = 216, N = 40) = .9839.$$

To make a reasonable allocation of subjects with unknown training status to cell 1 or cell 2, the predictive probability distribution should be followed. A split of 34 subjects to cell 1 and 6 subjects to cell 2 is most probable. A split of 33 subjects in cell 1 and 7 subjects in cell 2 is next most probable, and so on.

Exactly the same reasoning applies to the predicted fail column where we would like to make a reasonable allocation of the 156 students with an unknown training status. The predictive distribution for this allocation problem is presented in Fig. 6. Because N is larger (156 vs. 40) and because the known frequencies are smaller for cells 3 and 4, the distribution is more spread out. Only the range from $z = 80$ to $z = 140$ (where most of the probability lies) is shown on the graph. That most of the predictive probability distribution is contained within this range can be confirmed by the calculation

$$P(80 \leq z \leq 140|n = 63, y = 44, N = 156) = .9944.$$

For this case, what constitutes a reasonable split is more uncertain. The most probable value of z occurs at $z = 109$ with

$$P(z = 109|n = 63, y = 44, N = 156) = .0375$$




Figure 6: *The Bayesian predictive distribution for the allocation of 156 subjects who were predicted fails into the predicted fail-actual pass cell.*

but the graph shows a smooth progression with only small incremental changes in probability as we step forward and backwards from the mode. For example,

$$P(z = 108|n = 63, y = 44, N = 156) = .0373$$

and

$$P(z = 110|n = 63, y = 44, N = 156) = .0374.$$

As a result, there is a broader range of splits that are reasonably probable than in the first allocation problem. As seen earlier, 109 students allocated to cell 3 and 47 students allocated to cell 4 is most probable, but 100 students in cell 3 and 56 students in cell 4, or 120 students in cell 3 and 36 students in cell 4 are not unreasonable allocations either.

DETERMINING AVERAGE NET SAVINGS

Up to this point, we have examined these two allocation problems separately. But, in fact, to determine Δ , the difference in attrition rates, both allocations must occur together. Fortunately, the two allocation problems are independent of each other. One concerns the subjects who are predicted passes while the other concerns subjects who are predicted fails. Since independence between the two predictive probability distributions holds, these probabilities can be multiplied to find their joint occurrence needed in order to calculate Δ . In essence, the next step is to calculate a probability distribution for the savings due to implementation of the LCAC Selection System. The savings distribution is a linear function of Δ , and Δ is a function of the independent probabilities of the two predictive distributions.

Any particular allocation strategy results in some Δ . For example, suppose we are interested in the probability of $\Delta = .0581$. This difference in attrition rates arises by allocating 34 students to cell 1, 6 students to cell 2, 109

students to cell 3, and 47 students to cell 4. This the most probable Δ because it is based on the two most probable allocations. See Fig. 7 for the 2×2 table reflecting the joint occurrence of these two independent allocations. The point estimate for the probability of failure without the selection system is

	Predicted Pass	Predicted Fail	
Actual Pass	216 + (34) 250	44 + (109) 153	403
Actual Fail	43 + (6) 49	19 + (47) 66	115
	299	219	518

Figure 7: The most probable joint allocation of 196 subjects whose training status is unknown to the four cells of the 2×2 table.

$$\begin{aligned}
 P(\text{Attrite without selection system}) &= \frac{115}{518} \\
 &= .2220.
 \end{aligned}$$

The point estimate for the probability of failure with the selection system in place is

$$\begin{aligned}
 P(\text{Attrite with selection system}) &= \frac{49}{299} \\
 &= .1639.
 \end{aligned}$$

Therefore, the difference in the two attrition rates is

$$\Delta = .2220 - .1639 = .0581.$$

In the previous section, the predictive probability for each of these allocations was found. By multiplying the predictive probabilities for these two independent events, the result is

$$\begin{aligned}
 P(z_{\text{Cell } 1} = 34 \text{ and } z_{\text{Cell } 3} = 109) &= .1560 \times .0375 \\
 &= .0059.
 \end{aligned}$$

A subscript with the appropriate cell number is used to identify the predicted frequency for each of the two allocation problems. This probability of .0059 is assigned to $\Delta = .0581$.

What we are really interested in is the distribution of the savings due to the various Δ s. For the one Δ just examined, this is

$$\begin{aligned}
 \text{Net savings} &= (.0581 \times 30 \times \$160,000) - \$125,000 \\
 &= \$153,880
 \end{aligned}$$

and this particular net savings has a probability of .0059. A computer program was written to take all the independent combinations for $z_{\text{Cell } 1}$ and $z_{\text{Cell } 3}$ and compute the net savings for each combination. That is, $z_{\text{Cell } 1}$ was decremented from 40 to 0 and within this loop $z_{\text{Cell } 3}$ was decremented from 156 to 0. The predictive probability was calculated for each $z_{\text{Cell } 1}$ and $z_{\text{Cell } 3}$. The probability of the joint occurrence of $z_{\text{Cell } 1}$ and $z_{\text{Cell } 3}$ was calculated from the individual predictive probabilities. Then, a Δ was calculated based on the particular values of $z_{\text{Cell } 1}$ and $z_{\text{Cell } 3}$, and the net savings was computed for this Δ . When these operations are taken over all combinations of allocations, a distribution of net savings results.

The final objective is to report the average net savings for this distribution and the standard deviation about this average. Consult Fig. 8 for what this kind of analysis arrives at as a justification for continued usage of the LCAC selection system. The average net savings is close to \$160,000 with a standard deviation of about \$100,000.

Average	=	\$158,236
Standard Deviation	=	\$99,254
-2 SD	=	-\$40,272
+2 SD	=	\$356,744

Figure 8: *The average net savings due to implementation of the LCAC selection system. If a normal curve is used to approximate the distribution, the bottom two rows bracket about 95% of the net savings distribution.*

Assuming the Gaussian distribution as an approximation to the net savings distribution, then a 95% confidence interval around the average is about $-2sd$ at the low end, and $+2sd$ at the high end. This roughly brackets the net savings due to implementation of the LCAC Selection System between no savings and \$350,000, as given in the Introduction.

It is important to note that *every* possible allocation strategy has been included in this average. Each allocation strategy has been weighted by the predictive probability distribution. Thus, more weight is given to the "reasonable" allocation strategies that follow the empirical data and less weight to the "extreme" allocation strategies that do not.

REFERENCES

1. Blower, D. J. *Monitoring Operational Selection Systems Through Frequency Counts: An Application of Bayesian Predictive Inference*. NAMRL-1405, Naval Aerospace Medical Research Laboratory, Pensacola, FL, 1999.
2. Blower, D. J. *Analysis of Rejection and Attrition Data from the Landing Craft Air Cushion (LCAC) Selection System*. NAMRL-SR99-1, Naval Aerospace Medical Research Laboratory, Pensacola, FL, 1999.
3. Blower, D. J. *The Prediction Algorithm for the Landing Craft Air Cushion (LCAC) Selection System*. NAMRL-1395, Naval Aerospace Medical Research Laboratory, Pensacola, FL, 1998.
4. Blower, D. J. *A Quantitative Framework for Conducting Cost Benefit Analyses of Selection Tests*. NAMRL-1397, Naval Aerospace Medical Research Laboratory, Pensacola, FL, 1997.

APPENDIX

THE COMBINATORIAL ARGUMENT

In the predictive probability distribution, we always find the combinatorial factor,

$$\binom{N}{z} = \frac{N!}{(N-z)! z!}.$$

It gets modified by other terms in the predictive distribution (see Equation (19) in Blower [4]), but it has an influence in biasing an allocation strategy towards an even split as opposed to an extreme split.

This appendix provides a tutorial on this combinatorial factor. It's main purpose is to provide an easy rationale for arguing against extreme splits in any allocation strategy. In addition, the ideas explained here also have a subtle, but profound, effect on all of scientific inference. The combinatorial factor underlies the concept of maximum entropy assignment to probability distributions. The interested reader may wish to consult Volume II of my *Introduction to Scientific Inference* for a thorough introduction to the maximum entropy principle.

How would you allocate the 156 subjects to cells 3 and 4 of the 2×2 matrix? Of course, there is no absolutely clear cut answer to this question, short of knowing the training outcomes for these subjects. But there is an argument that most people would accept as reasonable. This is the combinatorial argument and it goes as follows.

How many conceivable ways are there to divide up N candidates into two cells? Consider a small N so that the answer can be worked out easily enough. For example, if there are $N = 4$ candidates to be allocated to the predicted fail-actual pass cell and the predicted fail-actual fail cell, how many ways can this be done? There are only five possible strategies of allocating these four candidates to two cells. See Table 1 for a listing of these five strategies.

Table 1: A listing of the five possible strategies for allocating four subjects to two cells.

Strategy	Cell 3	Cell 4	Number of ways
1	4	0	1
2	0	4	1
3	3	1	4
4	1	3	4
5	2	2	6

The final column gives the number of ways each of these strategies can be accomplished. This number is calculated from the combinatorial formula given at the beginning of this appendix. It depends on the fact that each of the four candidates is an individual and distinct person. For ease of explanation, call these four candidates *Alice*, *Bob*, *Carl*, and *Dawn*, or *a*, *b*, *c*, *d* for short. Under strategies 1 and 2, there is only one way all four people can be placed in a cell. Under strategies 3 and 4, however, the strategy can be achieved in four different ways depending upon who goes into cell 3. For example, under strategy 4, one way is for Alice to go into cell 3 and Bob, Carl, and Dawn to go into cell 4. The second way is for Bob to go into cell 3 and Alice, Carl and Dawn to go into cell 4. Now you can discern the pattern and easily find that the third way is for Carl to go into cell 3 and Alice, Bob, and Dawn to go into cell 4, while the fourth and final way is for Dawn to go into cell 3 and Alice, Bob, and Carl to go into cell 4.

For the fifth strategy, which represents an even split between the two cells, there are six ways to accomplish the strategy. Table 2 lists each one of the six distinct ways of allocating two candidates to cell 3 and two candidates to cell 4 using the shorthand notation for the names.

Table 2: The six possible ways of executing the fifth strategy of allocating two candidates to cell 3 and two candidates to cell 4.

Way	Cell 3	Cell 4
1	ab	cd
2	ac	bd
3	ad	bc
4	bc	ad
5	bd	ac
6	cb	ab

The point of this example is to emphasize that a strategy implementing a roughly even split into two cells is more likely than one that puts extreme counts into the two cells. For just $N = 4$ candidates, the ratio is only 4:1 or 6:1. When N becomes large, however, it is overwhelmingly more likely for a roughly even split between the two cells as opposed to more extreme counts. This argument is based solely on the combinatorial formula and has nothing to do with the *chance*, θ , of a candidate being assigned to one of the cells. Although θ does eventually get woven into the predictive distribution, right now we are highlighting the role of the combinatorial formula.

For larger N , some numerical examples reveal the impact of the combinatorial argument. In our current problem, 156 candidates need to be allocated to cells 3 and 4. Another combinatorial formula indicates how many strategies exist (call this number K) for N candidates allocated to n cells.

$$\begin{aligned}
 K &= \frac{(N + n - 1)!}{N! (n - 1)!} \\
 &= \frac{(156 + 2 - 1)!}{156! (2 - 1)!} \\
 &= \frac{157!}{156! 1!} \\
 &= 157.
 \end{aligned}$$

There are always just $K = N + 1$ strategies for N candidates to go into $n = 2$ cells.

Two of these $K = 157$ allocation strategies fall into the extreme category as discussed in the text. The first extreme strategy is to allocate all 156 candidates to cell 3 (the strategy most unfavorable to the LCAC selection system), and the second extreme strategy is to allocate all 156 candidates to cell 4 (the strategy most favorable to the LCAC Selection System). From our discussion of the example above, the combinatorial formula (as well as our unaided common sense) says that there is only *one* way of accomplishing these two extreme strategies. It is also easy to see that the next most extreme strategy of placing 155 candidates in cell 3 and 1 candidate in cell 4 can be accomplished in

$$\binom{156}{1} = \frac{156!}{155! 1!} = 156 \text{ ways.}$$

Now compare these numbers with those attached to a fairly even split. Take the even split allocation first. The strategy of placing 78 candidates in cell 3 and 78 candidates in cell 4 can be accomplished in

$$\binom{156}{78} = \frac{156!}{78! 78!} = 5.83 \times 10^{45} \text{ ways.}$$

This kind of ratio comparing an even split to an extreme split is what is meant by the phrase “is overwhelmingly more likely.” Now, as mentioned above, this does not mean that the even split receives the highest probability in the predictive distribution. The predictive distribution also takes into account the actual empirical data, but the combinatorial formula greatly influences or modulates towards allocations that are evenly split. Thus, a split of 109 candidates in cell 3 and 47 candidates in cell 4 can be accomplished in

$$\binom{156}{109} = \frac{156!}{109! 47!} = 2.00 \times 10^{40} \text{ ways.}$$

When numbers of this sort are combined with the frequency counts from the observed data, this allocation becomes the most probable of all.

The great contribution of the Bayesian approach is to provide a quantitative way of combining information from the combinatorial formula and the observed data. We can see the vague outlines of how to do this intuitively with very small numbers, but our common sense fails us when we are forced to deal with large numbers. The predictive formula is just one example of the self-consistent, disciplined approach based on probability theory to matters of scientific inference.

Reviewed and approved 15 SEP 00



R. R. STANNY, Ph.D.
Technical Director



This research was sponsored by the Naval Sea Systems Command under grant number N0002400WR00683

The views expressed in this article are those of the author and do not necessarily reflect the official policy or position of the Department of the Navy, Department of Defense, nor the U.S. Government.

Reproduction in whole or in part is permitted for any purpose of the United States Government.

REPORT DOCUMENTATION PAGE			Form Approved OMB No. 0704-0188	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.				
1. AGENCY USE ONLY (Leave blank)		2. REPORT DATE 15 SEP 00		3. REPORT TYPE AND DATES COVERED
4. TITLE AND SUBTITLE A Cost-Benefit Analysis of the Landing Craft Air Cushion (LCAC) Selection System			5. FUNDING NUMBERS N0002400 WR00683	
6. AUTHOR(S) D. J. Blower				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Naval Aerospace Medical Research Laboratory 51 Hovey Road Pensacola Fl 32508-1046			8. PERFORMING ORGANIZATION REPORT NUMBER NAMRL Special Report 00-2	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) Naval Sea System Command PMS 377C 2531 Jefferson Davis Hwy Arlington, VA 22242-5160			10. SPONSORING / MONITORING AGENCY REPORT NUMBER	
11. SUPPLEMENTARY NOTES Approved for public release; distribution unlimited.				
12a. DISTRIBUTION / AVAILABILITY STATEMENT			12b. DISTRIBUTION CODE	
13. ABSTRACT (Maximum 200 words) The purpose of this report is to document a cost-benefit analysis of the LCAC Selection System for Craftmasters and Engineers. A cost-benefit analysis for this selection system had not been conducted before and it seemed worthwhile to see if there was a cost justification for the continued use of this system. The analysis in this paper indicates an annual net savings somewhere in the range of no savings to \$350,000. The best guess is an annual net savings of about \$160,000. About 70% of the distribution is centered on the range of \$60,000 to \$260,000 net savings per year. Because the bulk of the distribution covers an expected cost benefit to the LCAC training commands, we recommend the continued usage of the LCAC Selection System to pre-filter candidates for training as Craftmasters and Engineers. Monitoring of the data and updates to the cost structure should be carried out periodically to determine if these savings can be expected to continue into the future.				
14. SUBJECT TERMS Personnel selection, Cost-benefit analysis, Bayesian statistics, Predictive distribution			15. NUMBER OF PAGES 24	
			16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT UNCLASSIFIED	18. SECURITY CLASSIFICATION OF THIS PAGE UNCLASSIFIED	19. SECURITY CLASSIFICATION OF ABSTRACT UNCLASSIFIED	20. LIMITATION OF ABSTRACT SAR	