



AFTAC

Air Force Technical Applications Center

Directorate of Nuclear Treaty Monitoring


**An Action-Space/Expected-Cost-of-Classification (ECC)
Approach to Theoretical Seismic Discrimination: Undecided
Regions, Unequal Population Variances, Costs and Benefits, Prior
Probability, Outlier Analysis, Three or More Populations, Test
Sites, and Variation of Discrimination Threshold with Magnitude**

Robert R. Blandford

2 April 2002

**Approved for Public Release;
Distribution is Unlimited**


20031217 225



Report AFTAC-TR-02-004 has been reviewed and is approved for publication.



DAVID R. RUSSELL, SES
Director, Nuclear Treaty Monitoring



ROY E. HORTON III, Colonel, USAF
Commander

The overall classification of this report is UNCLASSIFIED

Addressees: Please notify AFTAC/TT, 1030 S. Highway A1A, Patrick Air Force Base FL 32925-3002, if there is a change in your mailing address (including an individual no longer employed by your organization) or if your organization no longer wishes to be included in the distribution of future reports of this nature.

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 this 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 Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. **PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.**

1. REPORT DATE (DD-MM-YYYY) 2 April 2002		2. REPORT TYPE Technical		3. DATES COVERED (From - To)	
4. TITLE AND SUBTITLE A Decision-Space/Expected-Cost-of-Classification (ECC) Approach to Theoretical Seismic Discrimination: Undecided Regions, Unequal Population Variances, Costs and Benefits, Prior Probability, Outlier Analysis, Three or More Populations, Test Sites, and Variation of Discrimination Threshold with Magnitude				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
				5d. PROJECT NUMBER	
6. AUTHORS Robert R. Blandford				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Air Force Technical Applications Center (AFTAC/TT) 1030 S. Highway A1A Patrick AFB FL 32925-3002				8. PERFORMING ORGANIZATION REPORT NUMBER AFTAC-TR-02-004	
9. SPONSORING/MONITORING AGENCY				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBERS	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for Public Release; Distribution Unlimited					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT Some procedures for discriminating between earthquakes (Q) and explosions (X) set aside a region of the discrimination parameter space (x) in which no decision is made; this may be called the "unidentified" or "undecided" region, (U). The existing statistical literature seems not to explicitly provide any such option, although an undecided region may be rigorously supported by the decision theory literature on "action spaces." In this report we show how the concept of U arises naturally from the concept of the costs of classification: positive costs from misclassifications and "no decisions", and negative costs (benefits) from correct classifications. The resulting approach is a generalization of the "Expected Cost of Misclassification" (ECM) approach; we call the generalization the "Expected Cost of Classification" (ECC) approach. We also show how thresholds for detecting X as outliers of a Q population may be derived from cost considerations together with uniform distributions for X, and, together with plausible prior probabilities for Q and X, lead to reasonable thresholds in realistic scenarios. (Continued on Next Page)					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
					Robert R. Blandford
a. REPORT	b. ABSTRACT	c. THIS PAGE			19b. TELEPHONE NUMBER (include area code)
UNCLASSIFIED	UNCLASSIFIED	UNCLASSIFIED	UNCLASSIFIED	50	(703) 648-6789

Block 14 Continued:

Some procedures for discriminating between earthquakes (Q) and explosions (X) set aside a region of the discrimination parameter space (x) in which no decision is made; this may be called the "unidentified" or "undecided" region, (U). The existing statistical literature seems not to explicitly provide any such option, although an undecided region may be rigorously supported by the decision theory literature on "action spaces." In this report we show how the concept of U arises naturally from the concept of the costs of classification: positive costs from misclassifications and "no decisions", and negative costs (benefits) from correct classifications. The resulting approach is a generalization of the "Expected Cost of Misclassification" (ECM) approach; we call the generalization the "Expected Cost of Classification" (ECC) approach.

We also show how thresholds for detecting X as outliers of a Q population may be derived from cost considerations together with uniform distributions for X, and, together with plausible prior probabilities for Q and X, lead to reasonable thresholds in realistic scenarios. (Continued on Next Page)

Block 13 Continued:

Existing procedures also are sometimes one-sided, theoretically giving decisions only if the event is actually Q. For example, if an event is deep, it is Q; but if shallow, as is required for X and is possible for Q, there is no decision. We show how a similar concept can emerge naturally from the concept of different population variances for Q and X; the linear discriminant no longer applies, and the problem lies in the domain of quadratic discriminants such that there may be unconnected Q regions, X regions, and U regions. One also finds that one may have Q and U regions, but no X region.

In the fundamental approach to classification taken in this report, inversions of the estimated population covariance matrices are not required, whereas they are in classical linear discrimination. The more direct ECC approach simply involves estimation of the population descriptive parameters, such as means and variances, followed by direct determination, at each point of the decision space, of that population which has the lowest point ECC. The ECC is a function of the population means and variances, misclassification costs, "no decision" costs, the (possibly negative) costs (benefits) of correct classification, and the prior probabilities of Q and X.

Application of the method to a set of plausible International Monitoring Community parameters results in 80-90% probability of detection, one false alarm per year, and with a total cost 30-40% greater than the total benefits. For a system with better discrimination power, benefits would more nearly equal costs.

CONTENTS

	<u>Page</u>
List of Figures	vi
List of Tables	vii
Acknowledgments	viii
Abstract	1
Introduction	3
Cases of Interest	5
Theory: Expected Cost of Classification (ECC)	6
Equal Variance for Q and X	9
A Realistic International Monitoring Community Example	12
Test Sites	14
Unequal Variance for Q and X	15
Multiple Populations Example: Q and X Plus Mining (M)	16
An Analogy to Outlier Analysis: Uniform Priors for X:	22
Monitoring Realistic Seismicity: Variable Thresholds as a Function of mb	25
Summary and Suggestions for Further Work	28
Calculation Routines	29
References	30
Figures	31
Outline of Expected Cost of Classification (ECC) Theory	36
Summary	37
Distribution	38

LIST OF FIGURES

	<u>Page</u>
Figure 1	Typical hypothetical probability distributions 31
Figure 2	Plots of point costs and the ID parameter as a function of the discrimination variable, x 32
Figure 3	Results for unequal variance 33
Figure 4	Results for discrimination of Q, M, and X 34
Figure 5	Results for discrimination of Q and X with a uniform prior distribution for X 35

LIST OF TABLES

		<u>Page</u>
Table 1	Effects of Negative Costs (Benefits), and Unequal Priors and Variances	10
Table 2	Possible Cost (\$M) Matrix for Realistic Case	13
Table 3	Decision Costs for Mining/Seismic Region	16
Table 4	Mine Monitoring Scenarios of Interest	19
Table 5	Uniform Prior for X/Outlier Analysis	22
Table 6	Decision Costs as a Function of Magnitude ($p_x=1$)	26
Table 7	System Thresholds and Costs as a Function of Magnitude	27

ACKNOWLEDGMENT

Fruitful discussions with Bob Shumway are greatly appreciated.

ABSTRACT

Some procedures for discriminating between earthquakes (Q) and explosions (X) set aside a region of the discrimination parameter space (x) in which no decision is made; this may be called the "unidentified" or "undecided" region, (U). The existing statistical literature seems not to explicitly provide any such option, although an undecided region may be rigorously supported by the decision theory literature on "action spaces." In this report we show how the concept of U arises naturally from the concept of the costs of classification: positive costs from misclassifications and "no decisions", and negative costs (benefits) from correct classifications. The resulting approach is a generalization of the "Expected Cost of Misclassification" (ECM) approach; we call the generalization the "Expected Cost of Classification" (ECC) approach.

We also show how thresholds for detecting X as outliers of a Q population may be derived from cost considerations together with uniform distributions for X, and, together with plausible prior probabilities for Q and X, lead to reasonable thresholds in realistic scenarios.

Existing procedures also are sometimes one-sided, theoretically giving decisions only if the event is actually Q. For example, if an event is deep, it is Q; but if shallow, as is required for X and is possible for Q, there is no decision. We show how a similar concept can emerge naturally from the concept of different population variances for Q and X; the linear discriminant no longer applies, and the problem lies in the domain of quadratic discriminants such that there may be unconnected Q regions, X regions, and U regions. One also finds that one may have Q and U regions, but no X region.

In the fundamental approach to classification taken in this report, inversions of the estimated population covariance matrices are not required, whereas they are in classical linear discrimination. The more direct ECC approach simply involves estimation of the population descriptive parameters, such as means and variances, followed by direct determination, at each point of the decision space, of that population which has the lowest point ECC. The ECC is a function of the population means and variances, misclassification costs, "no decision" costs, the (possibly negative) costs (benefits) of correct classification, and the prior probabilities of Q and X.

Application of the method to a set of plausible International Monitoring Community parameters results in 80-90% probability of detection, one false alarm per year, and with a total cost 30-40% greater than the total benefits. For a system with better discrimination power, benefits would more nearly equal costs.

INTRODUCTION

All discrimination studies of which I am aware in the seismic literature, *e.g.*, Taylor *et al.* (1989), rely on linear discrimination. Linear discrimination assumes that the variance matrices of the discriminants are equal for Q and X, and that the costs for correct classification are zero. In addition, most studies implicitly assume that the prior probabilities of Q and X are equal; that the misclassification costs, $c(Q|X)$ and $c(X|Q)$, are equal; and that the costs (benefits) for correct classification are zero. For the theory of discrimination used in this memorandum, see Chapter 11 of Johnson and Wichern (1998).

However, some discriminants as implemented in this way may, in actual application, have a high false alarm rate, *e.g.*, 10% of Q may be classified as X, or 10% of X may be classified as Q. In addition to political costs, such events may result in costs for extensive additional technical work when it later becomes clear that a mistake was made. Thus, it is plausible that it would be useful to classify an event as "not discriminated" or "unidentified" if there were not high confidence that the event is Q or X, and if the costs of a failure to identify were less than that of an incorrect identification.

Of course, there is some cost, $c(U|X)$ or $c(U|Q)$, for each "no decision" in that further work, using other methods, must immediately be put into many such events. (In this report we use the notations $c(A|B)$ and "cab" interchangeably.) There would, of course, be less work put into a small, unidentified event in an uninteresting region than into a large unidentified event near a test site. We shall see that these differences can result in different decision and "no decision" thresholds as a function of event magnitude if the objective is to minimize cost.

In addition, there may be negative costs or "benefits", $c(X|X)$ and $c(Q|Q)$. It seems plausible that there would usually be a much higher benefit (more negative cost) for correctly identifying X than for correctly identifying Q.

Also, the costs of the overall discrimination process must be allocated in some manner, and it is plausible to do so by allocating a small cost to every event analyzed. The dominant positive costs, summed over all events will, in most cases, be $c(Q|Q)$ and $c(U|Q)$ because of the large number of Q events. In a "worthwhile" system one would expect the benefits, *e.g.*, $c(X|X)$ integrated over all X, to approximately cancel the total positive costs.

It is important to note that proof of the classical result that classification regions are determined by the maximum prior-weighted probability, an intuitively appealing result (*e.g.*, see Johnson and Wichern, 1998), depends not only on the assumption of equal misclassification costs and on zero values for $c(Q|Q)$ and $c(X|X)$, but also on the assumption that no other actions, such as U , are possible. If all of these assumptions are not suitable to the case of interest, then many classical discrimination approaches used in the literature will not be optimal with respect to cost.

If policy experts feel that the expected total political cost of a system's poor performance is too high, they can provide greater system resources, which should improve discrimination and may lower overall costs since, although the system costs will likely increase, the political costs may decrease even more.

The initial analysis cost for each event, which is the cost of the overall discrimination system operation as allocated equally to each event, can be seen in the mathematics not to affect the discrimination thresholds for a fixed, particular system. From a larger point of view, however, increased system costs should result in a different and better system with respect to discrimination, *e.g.*, there may be larger differences between the means or smaller population variances. Thus, the thresholds may change as system costs increase because the system itself changes.

The equally allocated cost does affect the total cost and, hence, can be an element to consider in choosing between different systems.

A number of studies of regional P/S discriminants show considerable overlap, especially at frequencies below 5 Hertz (Hz). However, it is often the case that a few Q may be identified by their especially low values. (See, *e.g.*, Figure 7 in Taylor *et al.*, 1989, who remark that the Q population shows substantially larger "scatter" than the X population.) This is reasonable on various physical grounds, and may be modeled by assuming that the Q variance is higher than the X variance. We shall see that this implies that linear discrimination does not apply, that some Q may be identified, but not X , and that there is a large unidentified (U) region of discrimination space.

Although we discuss only a single discrimination variable in this memorandum, it can be the case that multiple variables can be combined into a single variable. For example, Fisk *et al.*

(1995), Bottone *et al.* (1996), and Murphy *et al.* (1997) have shown how M_S , m_b , and the standard errors of both variables can be combined into a single variable. This variable could be treated by the techniques of this report.

Also, we may note that there is nothing intrinsically one-dimensional to the approach in this report; the probability densities, not necessarily normal, may be evaluated in two, three or more dimensions, and the same basic techniques utilized.

CASES OF INTEREST

Figure 1 shows the population probability distributions for several cases analyzed in this report. Figure 1a illustrates the case of a rather poor discriminant in which both variances are equal. The discriminant is poor because the standard errors are not small compared to the difference between the means.

The figure shows two normal distributions, $p(x)$, with means of -1 and 1 and a common variance of 1, representing earthquakes (Q) and explosions (X). To help fix ideas, the horizontal axis, x , may be thought of as a $\log(P/S)$ ratio which is smaller for Q than for X.

The standard linear discriminant, assuming equal prior probability, p , and equal costs of misclassification, c , would classify all events for which $x > 0$ as X and for $x < 0$ as Q. The false alarm rate for X, $P(X|Q)$, would be the cumulative normal, $\Phi(\mu, \sigma, x) = \Phi(-1, 1, 0) = 0.16$; and the probability of detection for X, $P(X|X)$, would be $1 - \Phi(1, 1, 0) = 0.84$.

Figure 1b illustrates the case of a good discriminant in which both variances are equal. The discriminant is good because the standard errors, 0.5, are small compared to the difference between the means, 2.0. In this case, the false alarm rate would be 0.02 and the probability of detection 0.98.

Figure 1c illustrates the case in which the variances are unequal: 1.0 and 0.5. In this case, the standard error of Q is twice that of X and the linear discriminant cannot be applied. We shall analyze this situation further below.

Figure 1d illustrates the case where there are three populations. To fix ideas we may imagine that the middle population represents mining explosions (M) which may have discrimination values

generally intermediate between those of Q and X. Here we have chosen to show all three populations with small standard deviations, 0.5. However, it is clear that the M population of events complicates the problem of identifying explosions. Plus, an entire new family of costs must be specified, *i.e.*, $c(X|M)$, $c(M|X)$, $c(U|M)$, $c(M|M)$, $c(M|Q)$, $c(Q|M)$. Presumably these last three or four costs would be small, compared to the first two.

In each of the cases in Figure 1, the X population could be replaced by a uniform distribution over some range. This might reflect a case where there have been no known nuclear explosions in a region of interest and so we imagine that we have little idea what values of x would be appropriate for an explosion, and so assume a uniform distribution over a large region of x . We shall show that in that case, analysis produces results very analogous to those obtained by outlier analysis (Taylor and Hartse, 1997; Fisk *et al.*, 1996) in which only the false alarm rate is controlled. Then, going further, we can show explicitly what the true costs and probabilities of detection appropriate for that false alarm rate would be if the true X distribution over x was some specific function.

In subsequent sections we shall take up these various cases of interest and examine how the thresholds and undecided regions change as the costs and priors are varied. We will also examine a set of costs and priors which seem plausibly realistic and observe how, as a result, thresholds could plausibly change as a function of magnitude, and how they might change near test sites and mines where priors change from those of a worldwide average.

But first it is necessary to outline the theory for the Expected Cost of Classification (ECC) method.

THEORY: EXPECTED COST OF CLASSIFICATION (ECC)

In order to motivate the theory, let us return to the case discussed in Figure 1a. A user, applying this discriminant, would likely find the costs of 16% false alarms unacceptable and, if he applied the discriminant at all, would actually decide that an event was a Q or X only if it had a small or large value, respectively, of x . Thus, implicitly, the analyst would be implementing an undecided region.

One formal approach to this problem would be to first find a value of x for which $P(X|Q)$ is an acceptable value, *e.g.*, 0.01. This value for x is $x=1.33$; *i.e.*, $1-\phi(-1,1,1.33)=0.01$, where ϕ is the

cumulative normal distribution. Thus, an explosion is declared only if $x > 1.33$. Then we find the value of x for which $P(Q|X)$ is < 0.01 , *i.e.*, -1.33 . The interval $(-1.33, +1.33)$ would comprise an unidentified (U) region. $P(X|X)$ under this procedure is $1 - \Phi(1, 1, 1.33) = 0.37$. The identical numbers apply to Q.

So, only 0.37 of events are identified, instead of the 0.84 with the standard linear discriminant, but f has been reduced to 0.01 from 0.16. The obvious question is: what are the relative costs of false alarms and unidentified events?

Johnson and Wichern (1998) give the fundamental criterion for discrimination of multiple populations from which special cases may be derived for equal covariance, equal prior probability, equal misclassification costs, zero benefits for correct classification, and only two populations. (The case usually treated typically makes all these assumptions.)

Theory (*e.g.*, see Ferguson, 1967) shows that classification may be generalized through the concept of “action spaces”. In this generalization, the decisions to be made are not explicitly the identity of the true event type given x , but instead the action to be taken, given x .

With this generalization, the fundamental result is that the optimal classification procedure amounts to choosing mutually exclusive and exhaustive action/classification regions, R_i , such that the prior-probability-weighted expected cost of classification of all populations is minimized.

It may be proved that this may be done by allocating the event generating x to the k th action for which the point cost, C_k ,

$$C_k(x) = \sum_i p_i \cdot f(x)_i \cdot c(k|i) \quad (1)$$

is smallest. The range (not necessarily contiguous) of x over which the k th expression is the minimum is R_k . Similarly, a particular x may be determined to be in that R_k for which the k th expression is the minimum over all k expressions. (Since $c(i|i)$ is often assumed to be zero, (1) is often expressed as a sum for which i is not equal to k . We shall discuss this aspect further below.)

Here f_i is the probability density distribution of population i ; c is the cost if the event generating x is in population i and results in action k ; and p_i is the prior probability of an event being in population i .

In (1), i represents true event type, *e.g.*,

explosion	X
earthquake	Q
mine blast	M

and k represents a possible action, *e.g.*,

decide	X
decide	Q
decide	M
decide	U (unidentified)

To further fix thoughts let us consider again Figure 1a in which we see two normal distributions with equal variance of 1.0 and means of -1.0 for Q and +1.0 for X. Thus, we have Q and X true event types; let us assume that the actions are "decide Q", "decide X", and U, and that we set the costs as $c(Q|X)=c(X|Q)=0.5$, and $c(U|X)=c(U|Q)=0.1$. It is important to note that we have set the costs for U to be less than the costs for misidentification. In this case we also set $c(Q|Q)=c(X|X)=0$.

Note that, in view of $c(Q|Q)=c(X|X)=0$, the only density function in C_Q is f_X , and in C_X , f_Q .

To the left of $x=0$ in Figure 1a, the smallest point cost will be C_Q , since there $f_X(x)$ is small and C_Q contains no $f_Q(x)$, which would be large. Therefore, to the far left we have R_Q , as is appropriate since the Q population dominates for negative values of x . Similarly, to the far right we have R_X . Symmetrically in-between we would have R_U , if it exists. (There is no R_U if all the misidentification costs are equal, because in that case C_U is the sum of C_Q and C_X , which are both positive in this case, so that there is no x for which C_U is the minimum.)

Note that, in this analysis, no use has been made of the assumption that the distributions are normal, have equal variances, *etc.*, and, in fact, the analysis is completely general for any distribution. We shall see that more complex sets of distributions simply result in more complex decision regions.

EQUAL VARIANCE FOR Q AND X

Using the decision cost parameters, 0.5 and 0.1, discussed above, together with equal priors, $p_X=p_Q=0.5$, we may plot the point cost expressions (1) as seen in Figure 2a. We find that $C_Q=C_U$ at $x_L=-0.70$. By symmetry the undecided region is between -0.70 and +0.70.

The total expected cost for one event (note $p_X + p_Q = 1.0$) is $c=0.056$, calculated as the area beneath the lowest point cost curves. It may be explicitly calculated that this cost is less than the cost would be if there were no undecided region or, indeed, for any other smaller or larger undecided region. Thus, as the theory guarantees, we have found the minimum cost set of regions. The false alarm rate (f) is 0.044 and the probability of detection (d) is 0.6 (compare to 0.16 and 0.84 for the linear discriminant).

The f and d values for the minimum cost solution are the probability that a single member of Q will be a false alarm, and that a single member of X will be detected, given the calculated thresholds. They are not weighted by priors, and cost coefficients have no influence, given the calculated thresholds, which were, of course, calculated using cost information. With minimum expected classification cost as the optimality criterion, values of f and d are no longer directly determinative optimality measures, although they are, of course, of interest and should have "reasonable" values.

If we consider the case where $c(U|X)=c(U|Q)=0.5$, the same as $c(Q|X)$ and $c(X|Q)$, then there is no undecided region and the cost is calculated to be 0.08. As expected, this cost is greater than with the optimally determined undecided region.

The two cases discussed above are tabulated in rows 1 and 2 of column 1 of Table 1.

It is clear that if the cost of an unidentified event is small, then minimization of total cost will result in a large undecided region. As an example of this behavior we can reduce $c(U|Q)=c(U|X)$ to 0.02 and find that the undecided region is now $[-1.6, +1.6]$, and that the total cost, c , has been reduced to 0.017 from 0.056. Of course, while f is now only 0.0046, d is reduced to only 0.27 instead of being 0.62 or 0.84.

In the next section we introduce the concept of a negative cost, or benefit, $c(X|X)=-0.5$, and examine how this changes the respective thresholds. The results of the calculations are given in Figure 2b and Table 1 (row 3, column 1).

Table 1: Effects of Negative Costs (Benefits), and Unequal Priors and Variances

Fig 2,3	For each case: $\mu_Q=-1.0$; $\mu_X=+1.0$; $c_{QX}=c_{XQ}=0.5$; $c_{QQ}=0.0$		
	$\text{sig}_Q=1.0$, $\text{sig}_X=1.0$ (Figure 2)	$\text{sig}_Q=.5$, $\text{sig}_X=.5$	$\text{sig}_Q=1.0$, $\text{sig}_X=.5$ (Figure 3)
	$c_{uQ}=.5$, $c_{uX}=.5$, $c_{XX}=0$, $p_X=.5$, $p_Q=.5$ Q:0.0:X $c=.080$, $f=.16$, $d=.84$	Q:0.0:X $c=.011$, $f=.02$, $d=.975$,	Q:0.18:X:3.18:Q $c=.042$, $f=.12$, $d=.95$
a	$c_{uQ}=.1$, $c_{uX}=.1$, $c_{XX}=0$, $p_X=.5$, $p_Q=.5$ Q:-0.7:U:+.7:X $c=.056$, $f=.044$, $d=.62$	Q:-.16:U:.18:X $c=.0087$, $f=.009$, $d=.95$	Q:-.1:U:.52:X:2.82:U:3.46:Q $c=.033$, $f=.064$, $d=.83$
b	$c_{uQ}=.1$, $c_{uX}=.1$, $c_{XX}=-.5$, $p_X=.5$, $p_Q=.5$ Q:-.68:U:-.2:X $c=-.14$, $f=.21$, $d=.86$	Q:-.16:U:-.04:X $c=-.23$, $f=.027$, $d=.98$	Q:-.1:U:.1:X:3.26:U:3.46:Q $c=-.20$, $f=.14$, $d=.96$
c	$c_{uQ}=.1$, $c_{uX}=.1$, $c_{XX}=0$, $p_X=0.09$, $p_Q=0.91$ Q:.46:U:1.86:X $c=.025$, $f=.002$, $d=.19$	Q:.12:U:.48:X $c=.0045$, $f=.0015$, $d=.85$	Q:.4:U:2.96:Q $c=.02$, $f=.08$, $d=0$
d	$c_{uQ}=.5$, $c_{uX}=.5$, $c_{XX}=-.5$, $p_X=0.09$, $p_Q=0.91$ Q:.46:U:.96:X $c=.0073$, $f=.025$, $d=.52$	Q:.12:U:.24:X $c=-.037$, $f=.0065$, $d=.94$	Q:.4:U:.68:X:2.66:U:2.96:Q $c=-.0028$, $f=.046$, $d=.74$
Table Notes: As examples, Q:-.7:U:.7:X indicates that Q is decided for $x < -.7$, U is decided for $-.7 < x < 0.7$, and X is decided for $x > +.7$; c_{XQ} is the classification cost for X given Q; p_Q is the prior probability for Q; c is cost.; f is the false alarm rate due to a single member of Q; d is probability of detection of a single member of X. Columns 3 and 4 have the same values for c_{uQ} , c_{uX} , c_{XX} , p_X , and p_Q as does the corresponding row in column 2.			

It is interesting to note, in Figure 2b, that for some ranges of x some point costs from equations (2) are negative due to the existence of negative conditional costs (benefits). The least, including negative values, prior-weighted point cost determines the action or decision at each point in parameter space.

The thresholds, as seen in Figure 2b, are now asymmetrical, $[-.68, -.2]$, expanding R_X , instead of $[-.7, +.7]$ due to the fact that there is a benefit for correctly identifying X, but not for correctly identifying Q, ($c(Q|Q)=0$). There is still an undecided region, U, for the minimum cost solution.

In the above calculation we have assumed equal cost for classifying Q and X as U; *i.e.*, $c_{uq}=c_{ux}=0.1$. In general, we would expect a greater cost for failure to identify X than for failure to identify Q. Repeating the calculation for Figure 2b, but with $c_{ux}=0.25$ and $c_{uq}=0.05$, we find only a modest shift in the thresholds: Q:-.8:U:-.24:X instead of Q:-.68:U:-.2:X. This may be interpreted as follows: as the cost for failure to identify X becomes greater, the integral of the explosion density over the unidentified region is reduced; and as the cost for failure to identify Q becomes less, the integral of the earthquake density over the unidentified region is increased.

In the above calculation, we have retained the assumption of equal prior probability, while making the costs for correct identification of Q and X asymmetrical. Now let us investigate the case where the costs of correct identification are restored to zero, but the prior probability of Q is 10 times that of X.

Again, the least-cost thresholds are seen, in Figure 2c and row 4 of Table 1, to become asymmetrical, [0.46, 1.86], in this case, as would be intuitively expected, greatly expanding the range of R_Q , shifting the undecided region toward the explosion population, and reducing the range of R_X .

While the cost is minimized for Figure 2c, and the false alarm probability is low, the probability of detection is poor, 0.19. This is an unsatisfying result; results of this sort have very likely led many researchers and policy experts to discount the use of priors. It seems counterintuitive to make it difficult to identify an explosion, just because there are many more earthquakes.

The reason that such analyses give such unsatisfactory results is that they do not take into account the cost benefits of correctly identifying explosions, especially as compared to the cost benefits of correctly identifying earthquakes. As we saw above, such cost considerations shift the threshold toward identifying more explosions. Only a proper balance of costs and priors will yield a reasonable result.

It is likely such weaknesses in existing theory have lead many users to rely directly on their intuition in setting f and d thresholds directly, instead of performing a more fundamental analysis.

In Figure 2d, we combine the two effects of a higher prior for Q and higher benefit for identifying X , and find the thresholds at $[0.46, 0.96]$. We see that the range of x over which explosions are identified has increased substantially; all X with x greater than the explosion mean are now identified as explosions; the probability of detection has increased to 0.52.

The foregoing analyses, with standard errors of 1.0, are repeated in the corresponding rows of column 2 of Table 1 for standard errors of 0.5. Some results of interest are that the undecided zones are smaller, and the costs are lower.

It is natural that the costs are lower because there are fewer errors made. However, this result is somewhat misleading because it is to be expected that to attain a better discrimination capability (smaller standard error), more funding might have to be provided by, for example, building more seismic arrays or investing in better processing capability. These costs would have to be allocated to the overall system and would perhaps best be specified as an equal cost for every analyzed event.

In particular, a cost such as $c(Q|Q)=0.01$, when allocated equally to each event and then summed over 1000 Q events and one or two X events, would massively affect the total cost. This cost would not, however, change any thresholds, or f or d values, unless, along with the cost, the mean and variance parameters changed. Of course, if funds were invested in a system, one would expect that the discrimination parameters would improve and, in that case, f and d would also be expected to improve, and costs, other than system enhancement costs, would decline.

If the costs of false alarms were high enough, then the sum of the reduced cost of a reduced number of false alarms, plus increased costs of system enhancement to develop better discrimination, could decline. We will discuss this point in more detail below.

A REALISTIC INTERNATIONAL MONITORING COMMUNITY EXAMPLE

Let us consider the situation where we are monitoring a large land mass where approximately 1000 events of interest, mostly earthquakes, occur per year. We may imagine that the cost of the monitoring system is approximately \$10M, so that $c_{qq}=\$0.01M$ processing cost per event. From

a policy point of view we assume that there could be one explosion per year in the region of interest. If we knew that there were no events, we would not be maintaining the monitoring system. An alternative assumption might be one explosion every 10 years, or 0.1 explosions per year.

We imagine that we have a fairly good discriminant capability, one as seen in Figure 1b, where the standard errors are 0.5.

Since we may presume that the policy-maker's principle interest in maintaining the monitoring system is to detect explosions, we set the benefit from detecting an explosion at the total expense of monitoring, *i.e.*, $c_{xx} = -\$10M$. (We may neglect the \$0.01 processing cost for the explosion.)

The cost, mostly political and strategic, for misidentifying an X as a Q, c_{qx} , may be taken as being of the same order as c_{xx} . For this example, we take it to be either \$2M or \$10M.

The costs for identifying Q as X are mostly political and are usually corrected by further analysis; we choose $c_{xq} = \$1M$.

The cost of an event being unidentified, U, is lower than being misidentified, because it may be assumed that further analysis will lead to a correct decision; and because that uncertainty will help to properly hedge political or tactical/strategic decisions. Failure to identify Q is less serious than failure to identify X, and so we take $c_{ux} = \$1M$, $c_{uq} = 0.2M$.

The foregoing may be summarized in Table 2.

Table 2: Possible Cost (\$M) Matrix for Realistic Case

$pq=.999, px=.001, \mu_{ux}=1.0, \mu_{uq}=-1.0; \sigma_{ix}=0.5; \sigma_{iq}=0.5$				
		Action		
		Decide:X	Decide:Q	U
Event Type	X	$c_{xx} = -10$	$c_{qx} = 2, 10$	$c_{ux} = 1$
	Q	$c_{xq} = 1$	$c_{qq} = .01$	$c_{uq} = .2$

The results of these calculations are that for $c_{qx} = \$2M$, the threshold between Q and X is at $x = 0.56$, approximately 1 standard deviation to the left of the mean of the X population, Q:0.56:X.

For $cqx = \$10M$, the threshold shifts slightly and a U region appears: $Q:0.4:U:0.54:X$. We see that the cost data have determined that the threshold be set so that there is approximately a 80-90% probability of detecting the explosion. For both cases, the probability of a false alarm is approximately 0.001, corresponding to approximately one false alarm per year.

The expected total cost for 1000 such events at the determined thresholds is \$3.1M for $cqx = \$2M$ and \$4.3M for $cqx = \$10M$. Thus, we see that the benefits do not fully cancel the costs for this set of discriminants. A better set of discriminants, *e.g.*, smaller standard errors, would result in benefits more nearly cancelling costs.

TEST SITES

If one is concentrating on a test site, then one often has experience suggesting that the prior probability of X is substantially greater than Q. So, plausibly for such a case, we choose $px = 0.909$ and $pq = 0.0909$ (a 10:1 ratio). Using parameters similar to those in Table 1, we choose for a test site: $mux = 1$, $muq = -1$, $sigx = sigq = 0.5$. This corresponds to good discrimination.

Then we choose $cxq = cqx = 0.5$. These correspond to a high cost for misidentification, as in Table 1.

Then we choose $cuq = cux = 0.25$. These correspond to a fairly high cost for an unidentified event as compared to Table 1 because, at a test site, unidentified events are more serious since the presumption is that an event is a test.

As for some cases in Table 1, we choose $cxx = -0.5$, there being a substantial benefit for identification of an explosion. However, we do not choose a benefit equal to the total cost of the monitoring system because the test site represents only a small portion of total monitoring. For a test site, there is also substantial benefit for correct identification of Q, because it helps to defuse possible political conflicts; we choose $cqq = -0.25$.

With these parameters there is a simple Q:X decision point at $x = -0.32$, and, as it happens, no U. The decision point is further to the left, toward the Q population, than any similar decision point in Table 1; to attain minimum cost there is a tendency to lean toward deciding that an event

is X. Because of the benefits of a high probability of a correct identification of X, the cost is less than for any example in Table 1; the cost is -0.47.

For any particular test site, it is clear that parameters other than those chosen here may be more suitable, and the decision points for minimum cost may be elsewhere.

UNEQUAL VARIANCE FOR Q AND X

Column 3 of Table 1 is for the case of unequal variance for the Q and X populations. The subplots, Figures 3a-d, correspond to Figures 2a-d except for the difference in X variance; *i.e.*, the standard error for X is 0.5 instead of 1.0.

Inspection of Figure 3a and the second row, third column of Table 1, shows that instead of having sequential Q:U:X regions as x increases, we now have sequential Q:U:X:U:Q regions. We are in the realm of the quadratic discriminant where the region resulting in identical decisions is not necessarily a single connected space. Basically, this increase in number of regions occurs because, for large x , the larger value of the standard error for Q ensures that $p_Q > p_X$ even though the mean of the X population is greater than the mean of the Q population. As would be expected, since the smaller standard error for X results in better discrimination, the cost is less than the cost when the standard error was 1.0 for both Q and X.

Again we may note in Figure 3b that, as the benefit for detecting X is large, R_X increases in size as compared to Figure 3a.

And in Figure 3c, where the prior on Q has increased, R_Q increases as compared to Figure 3a. In this case, due to the larger standard error for Q, there is, in fact, no R_X . The only choice is between Q or U. The probability of detection, d , is 0.0. If the discriminant under consideration were the only discriminant available, then one would have to examine each U by other means if one were to have any hope of detecting an X.

Where there is both a large benefit for detecting X and a large prior for Q, we have a result in Figure 3d intermediate between Figures 3b and 3c, and there exists an R_X .

MULTIPLE POPULATIONS EXAMPLE: Q AND X PLUS MINING (M)

We now address the question of extending the ECC approach to more than two populations. We shall see that it is actually a simple matter: all that is required is to continue the approach of evaluating each point cost expression (1) and defining that range of x where the k th is a minimum as R_k .

We shall generalize the Q and X populations by adding a mining event population, M, with mean, 0.0, between those of the Q (-1.0) and X (+1.0) means. As noted previously, it is necessary to define a number of additional costs. Here, in Table 3, as an illustration we give detailed plausible estimates of these costs

It is useful, in thinking about these costs, to categorize them into monitoring and political costs and, within these categories, to further subdivide into short-term and long-term costs. The total cost for each conditional cost is the sum over short- and long-term costs, and over monitoring and political costs.

Consider, for example, the first row in Table 3 which gives values for c_{qx} ($c(Q|X)$): the cost for misclassifying an explosion as an earthquake. Since, in operations, it may not be known immediately that there has been a mistake, one may imagine that there is no short-term cost, either monitoring or political. The detection of an earthquake does not provoke any special studies in a monitoring system and the political system is indifferent. Thus, both short term entries for c_{qx} are set to 0.0 in Table 3.

Table 3: Decision Costs for Mining/Seismic Region

Cost Type	Monitoring System		Political System		Sum
	Short Term	Long Term	Short Term	Long Term	
c_{qx}	+0.0	+0.5	+0.0	+0.5	+1.0
c_{xq}	+0.1	+0.5	+0.1	+0.5	+1.2
c_{ux}	+0.1	+0.1	+0.0	+0.2	+0.4
c_{uq}	+0.1	+0.1	+0.0	+0.1	+0.3

(continued on next page)

Table 3: Decision Costs for Mining/Seismic Region (Continued)

Cost Type	Monitoring System		Political System		Sum
	Short Term	Long Term	Short Term	Long Term	
cxx	+0.1	+0.0	-0.2	-0.5	-0.6
cqq	+0.01	+0.0	+0.0	+0.0	+0.01
cqm	+0.1	+0.1	+0.0	+0.2	+0.4
cxm	+0.1	+0.25	+0.1	+0.25	+0.7
cum	+0.1	+0.1	+0.0	+0.2	+0.4
cmx	+0.0	+0.5	+0.0	+0.5	+1.0
cmq	+0.1	+0.1	+0.0	+0.2	+0.4
cmm	+0.01	+0.0	+0.0	+0.0	+0.01

On the other hand, in the long-term, there may be large political costs for having missed the detonation of a nuclear test; and once the fact that an apparent earthquake was actually a nuclear test is discovered, the monitoring system is likely to devote a substantial amount of resources to studying the situation and to implementing means to avoid a repetition of the incident.

Thus, it is reasonable that the long-term cqx costs would be nonzero for both monitoring and political systems; we set them equal to 0.5. We shall discuss in a subsequent section the relations of these cost numbers to plausibly realistic dollar budgets.

The second row, for cxq, is much the same except that there is immediate short-term work created for both the monitoring system and the political system when an explosion is thought to have been detected.

The values entered in the third row, cux, reflect the idea that having an explosion determined to be unidentified is not as costly as being identified as an earthquake, and also reflect the idea that there is some immediate work generated in the monitoring system when an event is unidentified.

As discussed previously, cuq is plausibly less costly than cux.

The large negative costs (benefits) in the political system for the fifth row of Table 3, cxx, reflect the idea that correct identification of explosions is the principle goal of the monitoring system and each such rare event represents a payoff which should be at least comparable to the total cost of the system.

For the sixth row of Table 3, cqq, we see that the political costs are 0.0. There is no political interest in the correct identification of earthquakes. However, from the monitoring point of view, the chief activity of the monitoring system is the correct identification of earthquakes so that one may be sure that they are not explosions. The total expense of the system should be allocated over all the events analyzed. Most of these events are earthquakes and so a small, short-term cost, 0.01, is allocated to such events. (This cost may actually be regarded as present as a short-term monitoring cost in each of the other categories so that, for example, the total 0.5 monitoring cost for cqx should be regarded as being made up of 0.01 routine costs and 0.49 extraordinary costs. The differences in thresholds resulting from a cost of 0.49 as compared to 0.50 are, of course, negligible.)

The small monitoring cost per event comes significantly into play only when (1) there are no other extraordinary costs, as in cqq and cmm; and (2) at the same time, there are a great number of such events. This combination is often the case for small earthquakes.

It is reasonable to assign substantial costs to cqm, cum, and cmq, because confusion as to whether an event is a mine event or not makes it difficult to see clearly what is happening and thus reduces confidence in the discrimination procedures.

Additional rationales for the costs in Table 3 should be apparent by analogy to the paragraphs above. More detailed discussions might only be useful in the context of an actual application.

With the set of costs from Table 3, Figures 4a-d and Table 4 give results for several cases of interest involving mining events.

Table 4: Mine Monitoring Scenarios of Interest

Fig	For each case: $\mu_q = -1.0$, $\mu_x = +1.0$, $\mu_m = 0$, $\sigma_q = 0.5$, $\sigma_m = 0.5$, $p_x = 1$, $p_m = 10$
4a	$\sigma_q = .5$, $p_q = 1$ Q:-1.06:M:.88:X $c = .71$, $f_q = .0001$, $f_m = .039$, $d = .59$
4b	$\sigma_q = 1$, $p_q = 1$ Q:-1.22:M:.88:X:3.2:U:3.32:Q $c = .71$, $f_q = .03$, $f_m = .039$, $d = .59$
4c	$\sigma_q = 1$, $p_q = 10$ Q:-.6:M:.94:U:1.02:X:2.56:U2.76:Q $c = 2.52$, $f_q = .02$, $f_m = .02$, $d = .48$
4d	$\sigma_q = .5$, $p_q = 10$ Q:-.48:M:.88:X $c = 1.75$, $f_q = .0001$, $f_m = .039$, $d = .59$
Table Notes: As examples, Q:-1.06:M:.88:X indicates that Q is identified for $x < -1.06$, M is identified for $-1.06 < x < 0.88$, and Q is identified for $x > +.88$; p_m is the prior probability for M; c is cost; f_q is false alarm rate due to Q; f_m is false alarm rate due to M; and d is probability of detection. All costs are from Table 3.	

In Table 4, the “priors” for X and M are taken as 1 and 10, respectively. Although true priors are correctly defined as probabilities which sum to 1.0 over all possibilities for a single event, in this example we multiply the true priors by the total number of events in the scenario. Using these “generalized” priors, the resulting cost, c , is then the expected cost for the scenario and not the expected cost for a single event. We shall see that this approach is useful when we examine the changes of thresholds and costs due to the increase in the number of events of interest as magnitude decreases.

As previously noted, we represent the population of mining explosions, M, by a normal distribution with mean, $\mu_m = 0$, midway between the populations of Q and X with means at -1 and 1, respectively. The standard errors for M and X are taken as 0.5, reflecting, generally speaking, good discrimination.

The first row in Table 4, with cost functions graphed in Figure 4a, may be viewed as a case wherein a particular mine or mining region is being monitored, and near which the relative activity of nuclear explosions and earthquakes are comparatively low as compared to the mining events. The discrimination capability for the earthquakes is taken also to be rather good, comparable to that of M and X; $\text{sigq}=0.5$.

We see that the cost is $c=0.71$, $f_q=.0001$, $f_m=.039$, and $d=0.59$, where f_q is the false alarm rate in which Q are misidentified as X, and f_m is the false alarm rate in which M is misidentified as X.

By comparison, if we set $p_m=0$, to revert to the simple Q and X case, but with the complex set of costs in Table 3, we find a single decision line such that $c=-0.52$, $f_q=0.02$, and $d=0.98$. Clearly, introducing the mining population has severely degraded discrimination and raised the cost, given that Q versus X discrimination capability remains the same, and that the same number of Q and X events are analyzed.

The second row in Table 4 (Figure 4b) shows that even if the standard error of Q is increased, $\text{sigq}=1.0$ (that is, discrimination capability for Q is poor), there is little effect on the cost. This is because the main effect of the increase in sigq is to make it more difficult to discriminate between M and Q. Since there are many more Q events, the decision point moves from $x=-1.06$ to -1.22 so that there are not too many cmq costs inflicted. Most of the cost is due to overlap between the M and X populations, and this is not changed by the increase in the standard error for Q. It is difficult to identify explosions at a mine.

The third row in Table 4 (Figure 4c) shows that if, in addition to increasing the variance of Q, we increase the priors so that there are equal numbers of Q and M, then there is a substantial effect on cost. This is plausible because one can no longer appeal to the greater number of M, as in the paragraph above, to allow shifting the decision line toward Q. So one must absorb the many cqm and cmq costs. Thus, a mining site in an active seismic area will be more costly to monitor than the same site in an aseismic region, as is plausible, given the specified costs.

Finally, if we attempt to improve the situation found in the previous paragraph by reducing the standard error of Q from 1.0 to 0.5, thus improving discrimination for Q, the fourth row in Table 4

and Figure 4d show that indeed the cost, c , is noticeably reduced; but we still cannot recover to the low costs of the low-seismicity regions, where $pq=1$, with $\text{sig}q=0.5$ or 1.0 ; the costs due to M can not be completely overcome.

Note that in all these scenarios the standard goodness parameters, f and d , remained close to the same values. Thus, we see that these parameters do not tell the whole story.

Note, however, that if in the case of Figure 4d, we reduce the costs of mistaking Q for M or M for Q (thus effectively to some degree lumping the two populations into one), the overall costs are reduced, and the thresholds are stable down to values of $c_{qm}=c_{mq}=0.01$. Such a set of costs would be suitable if one felt that there were a number of types of events and that it was not important to discriminate between them; it is only a requirement to be sure that they are not X . Such an approach has been discussed from the outlier point of view by Gray *et al.* (1996).

AN ANALOGY TO OUTLIER ANALYSIS: UNIFORM PRIORS FOR X:

In many areas of interest it may be that there are no sample events from the X population. In this case, we may want to make a less definite estimation of the distribution of X. Perhaps we could assume that the prior distribution of X is uniform over the range of the discrimination parameter.

First let us consider, in Figure 5a and in row 1 of Table 5, a set of parameters which we have considered previously, (Row 2, Table 1). We see that, as before, Q is identified on the left, U in $[-0.7, +0.7]$, and X to the right.

Table 5: Uniform Prior for X/Outlier Analysis

Fig	For each case: $\mu_q = -1.0$, $\sigma_q = 1.0$, $\mu_x = 1.0$, $p_x = 0.5$, $p_q = 0.5$, $c_{xq} = 0.5$, $c_{qx} = 0.5$, $c_{uq} = .1$, $c_{ux} = .1$, $c_{xx} = 0$, $c_{qq} = 0$
5a	$\sigma_x = 1.0$ Q: -0.7:U: +0.7:X $c = .056, f = .044, d = .62$
5b	X assumed uniform in $[-5, 5]$, $\sigma_x(\text{real}) = 1.0$ X: -3.34:U: 1.36:X $c = .077, cr = .085, f = .019, fr = .009, d = .53, dr = .36$
5c	X assumed uniform in $[0, 5]$, $\sigma_x(\text{real}) = 1.0$ Q: 0.04:U: 1.04:X $c = .023, cr = .071, f = .021, fr = .021, d = .40, dr = .48$
5d	X assumed uniform in $[-5, 5]$, $\sigma_x(\text{real}) = 0.5$ X: -3.34:U: 1.36:X $c = .077, cr = .092, f = .019, fr = .009, d = .53, dr = .23$
Table Notes: As examples, Q: -0.7:U: 0.7:X indicates that Q is identified for $x < -0.7$, event is U for $-0.7 < x < 0.7$, and Q is identified for $x > +0.7$; p_q is the prior probability for Q; C is cost; cr is real cost, with uniform prior for X assumed, if real distribution is normal; f is false alarm rate due to Q; d is probability of detection.	

Let us suppose, in Figure 5b and in row 2 of Table 5, that the X population is replaced by a uniform distribution in $[-5, 5]$. We then see that the U region is greatly enlarged, $[-3.34, 1.36]$, there is no Q region, and the X region exists both to the right and left. (The absence of a Q region,

as seen here, is not a general result. If $cuq > 0.1$, instead of $cuq=0.1$ in this case, then there would be a Q region symmetrical about $x=-1$. Then, as x further departed from -1.0 , there would be symmetrical U regions together with the remnants of the X regions seen in Row 2.)

Except for the existence of the U regions, this result seems very similar to those resulting from the outlier techniques developed by Fisk *et al.* (1996), and applied by Taylor and Hartse (1997). The present approach seems to add some flexibility to the outlier approach, making it possible to consider costs and priors; and to specify the region of discrimination space in which X may lie, if such information exists, instead of simply saying that it must lie outside some specified populations with some probability.

We can see, in Table 5, the kinds of qualitative behavior which one would expect when comparing the "outlier" approach to an approach where it could be assumed that there was some detailed knowledge of the X distribution.

Comparing row 1 to row 2 (Figures 5a and 5b) in Table 5, we see that as a result of the assumption of a uniform distribution for X, the overall cost has increased from $c=0.056$ to $c=0.077$. This latter cost is calculated on the basis of the assumption that the true distribution of X is uniform in $[-5, 5]$. If we assume that the true distribution for X is as seen in row 1, $N(1,1)$, the real cost, cr , for the threshold calculated assuming a uniform distribution, may be calculated (using Matlab routine `realcost.m`) to be $cr=0.085$. This is greater than the optimum cost which one would have with perfect information, and also more than the cost if a uniform distribution were true.

The outlier technique, because it makes no assumptions about the X population, controls only the false alarm rate, f . We see that, for the minimum cost solutions, f has decreased from .044 to .019 in transitioning from a known normal distribution to a uniform distribution, even though the cost has increased. (Note that for a uniform distribution, false alarms occur on both tails of the Q distribution.) But the probability of detection did decrease from $d=0.62$ to 0.53 , and this would lead to increased costs.

For the true distribution, f_r , the "real" false alarm rate, is calculated on only one tail of the Q distribution and $f_r=0.009$, and dr is less also, with a value $dr=0.36$.

The most important point in contrasting rows 1 and 2 of Table 5 is that the true cost increase, due to the assumption of a uniform distribution of X , is from 0.056 to 0.085. Of course, other sets of conditional costs than those given in Table 5 would result in different cost increases. Another way of looking at this result is that the outlier technique controls only the false alarm rate, it gives no prediction as to the probability of detection.

Row 3 in Table 5, and Figure 5c show how the uniform distribution could be restricted only to "reasonable" regions of x . In this case, by way of example, we assume a uniform distribution in $[0, 5]$. (Perhaps the region should extend a bit into negative x , considering the tails of the true X distribution, but we will pass over such details here.) Note that the resulting decision regions are much more similar to the "correct" decision lines seen in 5a than to the decision lines derived assuming a uniform distribution, as seen in 5b. We see that the costs decrease as compared to the previous case where there was a uniform distribution in $[-5, 5]$. A similar procedure might be useful in multivariate spaces if "reasonable" regions could be determined there also.

Row 4 in Table 5 (Figure 5d, corresponding to Row 4, is identical to Figure 5b) shows that if the real distribution for X is even narrower (more distinct from the uniform assumption of Row 2) then cr increases from 0.085 to 0.092.

In these analyses we have assumed that $c_{xx}=0$. If $c_{xx}<0$ so that there is a benefit to detecting X , then C_X decreases, and the unidentified region in Figure 5d decreases so that there are more X detected.

Consideration of these cases suggests that the ECC technique has the capability to reasonably handle cases when there is little detailed knowledge of the distribution for X , the situation for which the outlier technique was designed.

If we know more about the distribution of X , say that its population lies completely to the right of $x=0$, then we should be able to improve the costs. Row 3 (Figure 5c) as compared to row 2 of Table 5, shows that cr does decrease in this case from 0.085 for X in $[-5, 5]$ to 0.071 for X in $[0, 5]$.

Use of a uniform distribution could also be useful in multivariate discrimination where along some dimensions the discriminant is well determined for both Q and X, (e.g., $M_S:m_b$, and possibly depth), whereas along another dimensions (e.g., P/S) it is well determined for Q but not X. So, a uniform distribution would be assumed only for X along the P/S axis. Thus, the same basic discrimination processes can be implemented in all scenarios of interest; the only difference would be in the degree of specificity of the probability distributions.

MONITORING REALISTIC SEISMICITY: VARIABLE THRESHOLDS AS A FUNCTION OF m_b

As we have seen in the analyses above, thresholds that result in minimum cost will vary as a function of prior probability. Since the number of expected explosions may plausibly be regarded as fixed and small, while the number of earthquakes and mining blasts increases as magnitude decreases, it is plausible that thresholds for minimum cost monitoring will vary as a function of magnitude.

Other reasons for imagining that the thresholds would vary as a function of magnitude are that the efficiency of discrimination is likely to decline as a function of magnitude (the standard error is likely to become larger due to poorer S/N and fewer detecting stations); and the benefits and costs of correct and incorrect identification may decline in absolute value as magnitude decreases.

For the many small/unimportant events, we cannot afford a large expense per event to further process many unidentified events as compared to the costs for false alarms, i.e., $c(U|Q)$ may be very small compared to $c(Q|X)$ and $c(X|Q)$ for small magnitudes. In this case, to minimize total cost for such events, the unidentified region expands, the prior on X tends to dominate, and the proper strategy is to identify only those events for which the discriminant is unequivocal.

On the other hand, for the few large/important events, we can afford a greater expense to work on the unidentified events, i.e., $c(U|Q)$ may be larger compared to $c(Q|X)$ and $c(X|Q)$. In this case, the minimum total cost (including the costs of work on missed detections and false alarms) is found by using a smaller unidentified region.

It is worth noting that the absolute costs for large/important events are likely larger than for the small/unimportant events. This fact will come into play when one allocates a fixed budget between small and large event analysis, but it is only the ratio of costs which comes into play in determining the undecided zone for a fixed magnitude or other level of importance.

To begin to discuss this issue, Table 6 gives some figures which are meant to be illustrative costs which might be thought of as realistic in some situations.

Table 6: Decision Costs as a Function of Magnitude (px=1)

Cost Type	3<mb<4, pq=1000		4<mb<5, pq=100		5>mb, pq=10	
	Monitoring Short/Long	Political Short/Long	Monitoring Short/Long	Political Short/Long	Monitoring Short/Long	Political Short/Long
cqx	0.00/0.25	0.00/0.5	0.0/0.5	0.0/0.5	0.0/0.5	0.0/0.5
cxq	0.05/0.125	0.05/0.125	0.1/0.25	0.1/0.25	0.1/0.5	0.1/0.5
cux	0.1//0.1	0.0/0.1	0.1/0.1	0.0/0.1	0.1/0.1	0.0/0.2
cuq	0.05/0.0	0.0/0.0	0.05/0.05	0.0/0.05	0.1/0.1	0.0/0.1
cxx	0.1/0.0	-0.2/-0.5	0.1/0.0	-0.2/-0.5	0.1/0.0	-0.2/-0.5
cqq	(0.01/0.00)	0.0/0.0	(0.01/0.00)	0.0/0.0	(0.01/0.00)	0.0/0.0

Among the points of interest in Table 6:

- The prior probability for Q increases from 10 to 1000 as mb decreases.
- The costs of misidentification decline as mb decreases.
- The benefit of correct identification of X remains constant as mb decreases.
- The cost of misidentifying X as Q is 4-7 times greater than X as U.
- cqq is 0 or 0.01; if not negligible, then small compared to all other costs, so:
- Identifying Q as U is 5-10 times more costly than Q as Q; (If cqq not equal to 0.0).

The results of calculations given using the parameters of Table 6 are given in Table 7:

Table 7: System Thresholds and Costs as a Function of Magnitude

For each case: $\mu_q = -1.0$, $\mu_x = +1.0$, $\sigma_q = 0.5$, $\sigma_x = 0.5$, costs: Table 5		
$3 < mb < 4$, $pq = 1000$	$4 < mb < 5$, $pq = 100$	$5 < mb$, $pq = 10$
$Q:0.6:U:0.74:X$ $c = -0.123, f = 0.00025, d = 0.70$ $cS = 9.86$	$Q:0.4:U:0.5:X$ $c = 0.28, f = .0013, d = 0.84$ $cS = 0.72$	$Q:0.22:U:0.28:X$ $c = -0.42, f = 0.005, d = .925$ $cS = -0.32$
Table Notes: As examples, $Q:0.6:U:.74:X$ indicates that Q is identified for $x < 0.6$, U is identified for $0.6 < x < 0.74$, and X is identified for $x > 0.74$. pq is the prior probability for Q, c is cost if $cqq = 0$, and cS is the cost if $cqq = 0.01$. (Changing cqq from 0.00 to 0.01 changes thresholds only by approximately 0.01. These threshold changes can be neglected for all purposes in this memorandum.) f is the false alarm rate due to a single member of Q. d is probability of detection of a single member of X.		

We note that as the Q prior increases from 10 to 1000, the U region dividing Q and X moves to larger values of x , toward X, making it more difficult to identify X. Thus, the false alarm rate, f , and probability of detection decrease as magnitude decreases. Note, however, that the prior-probability weighted false alarm rate, $f_w = pq * f$, increases from 0.05 to 0.25. This means that in the course of 1000 low-magnitude Q events (a year of events), the expected number of Q that would be identified as X is 0.25.

We note that if $cqq = 0$, the cost of operation decreases as the magnitude decreases, probably due to the fact that the costs of misidentification decrease, while c_{xx} benefits are held constant. However, if we set $cqq = 0.01$ instead of 0.00, then the system cost increases as magnitude decreases. For example, for the lowest magnitude range, the cost is 9.86 instead of -0.123. For costs between 0.01 and 0.00, the system cost would be between 9.86 and -0.123.

For $cqq = 0.01$, the total cost, summing over all magnitudes, is 10.26. (Note that this is almost identical to 10.28, which is equal to 11.1, the cost for 1110 Q with a cqq of 0.01, plus -0.823, the total costs if $cqq = 0$. Thus, almost the total net effect of setting $cqq = 0.01$ is simply to add cqq for all Q.)

Thresholds were also determined using a uniform distribution for X; the major effect was to raise the thresholds in Table 6 by about 0.2.

Using these uniform distribution thresholds, costs were then estimated using the "true" distribution for X , and the cost was found to increase from 10.26 to 10.66. This small, 4%, increase in cost reflects the dominance of cqq costs over all others. If cqq costs were not so dominant, then the costs of misclassification would lead to a greater percentage difference in cost. For example, if $cqq=0$, then the change in true costs would increase approximately from -0.82 to -0.44, roughly a 50% increase in cost. So the benefit of knowing the "true" distribution of X varies between 4% and 50% of total system costs under these assumptions.

It is also worth noting that costs for the lowest magnitudes would greatly increase if the discrimination capability was substantially worse for the lowest magnitudes, as may often be the case in practice. For example, if $\text{sigx}=0.707$ for the lowest magnitude, and $cqq=0.0$, then the ECC cost, c , increases from -0.123 to 0.464, and pd decreases from 0.7 to 0.26.

SUMMARY AND SUGGESTIONS FOR FURTHER WORK

This memorandum has outlined a general procedure for discrimination which has most of the properties which experience has shown are desirable in practice. In addition, the decision thresholds emerge naturally from cost estimates which system managers are expected to be prepared, in practice, to make.

The procedure may be applied sequentially to several discriminants, and some sort of voting scheme applied to the output; this is similar to some present procedures. Alternatively, the procedure may be easily generalized to multiple dimensions.

In the former case, cost considerations might not easily be carried through to the final seismic decision; in the latter, multiple dimension case, they easily could be. Thus, a multiple dimension approach would appear to be more desirable.

However, a plausible approach to a sequential procedure would be to assume that a single Q or X identifies the event. Only if all discriminants decided U , would the event be U . Then, plausibly, all costs and benefits should be the same as in the multidimensional case except that all costs for identification of an event as U would be divided by the number of discriminants. In this way, if all discriminants resulted in U , then the calculated cost would be correct. It would be useful to make this argument rigorous.

The next project should be to apply this approach to some actual data resulting from operation of an operational system. Most useful, perhaps, would be Ms:mb and depth data, properly reflecting the detectability of seismicity, and also assuming a uniform distribution in magnitude for a small probability of explosions, also appropriately weighted by detectability. Another possible discriminant would be a phase ratio and a spectral ratio data from a monitoring site of interest.

Another project would be to consult with system managers and determine a definitive set of costs and benefits for use in actual applications.

CALCULATION ROUTINES

The actual calculations in this memorandum were performed by several Matlab routines written by the author. The routine `plfig.m` calculated the cost distributions for Q , U , and X , using equation (1), detected which was a minimum for each x and thus calculated the identification statistic (ID), printed out the transition values of x (which amounted to detection thresholds), and plotted the cost distributions and the ID statistic. Standard routines within `plfig.m` were used to calculate c , f , and d , given the thresholds.

The routine `plfig4.m` performed the same calculations as `plfig.m` except that the M population was added. Both routines had an option to assume a uniform prior distribution for X in $[-5, 5]$.

The routine `plfig_exp.m` performed the same calculations as `plfig.m` except that it modeled the X prior distribution as uniform between input parameters `xleft` and `xright`.

The routine `realcost.m` is used to calculate the cost for a set of thresholds and point cost distributions. The application is to use thresholds calculated assuming a uniform prior distribution for X , but assuming some other, true, normal distribution for X .

REFERENCES

- Bottone, S., M.D. Fisk, H.L. Gray, and G.D. McCartor (1996). The dependence of magnitude uncertainty on station coverage, PL-TR-96-2250, Phillips Laboratory, Hanscom AFB, MA.
- Ferguson, Thomas S. (1967). *Mathematical Statistics, A Decision Theoretic Approach*, Academic Press.
- Fisk, M.D., H.L. Gray, and G.D. McCartor (1995). Statistical methodology and assessment of seismic event characterization capability, PL-TR-95-2156, Phillips Laboratory, Hanscom AFB, MA, ADA305487.
- Fisk, M., H. Gray, and G. McCartor (1996). Regional event discrimination without transporting thresholds, BSSA, 86, 1545-1558.
- Gray, H.L., S.R. Sain, W.H. Frawley, and W.A. Woodward (1996). On detecting outliers in mixed populations, PL-TR-96-2049, Phillips Laboratory.
- Johnson, R.A., and D.W. Wichern (1998). *Applied Multivariate Statistical Analysis* (Fourth Edition), Prentice Hall.
- Murphy, J.R., B.W. Barker, and M.E. Marshall (1997). "Event Screening at the IDC Using the Ms/Mb Discriminant", Technical Memorandum to NTPO, October 1997.
- Taylor, S., M. Denny, E. Vergino, and R. Glaser (1989). Regional discrimination between NTS explosions and western United States earthquakes at regional distances, BSSA, 79, 1142-1176.
- Taylor, S., and H. Hartse (1997). An evaluation of generalized likelihood ratio outlier detection to identification of seismic events, BSSA, 87, 824-831.




Figure 1. Typical hypothetical probability distributions for populations of earthquakes (Q), explosions (X), and mining events (M) studied in this memorandum. The means of the distributions are centered either at -1, 0, or 1. The broader distributions have standard errors of 1.0, the narrower, 0.5. The general procedure in this memorandum does not depend on normality; normal populations are used only for convenience.




Figure 2. Plots of point costs (Equation 1) and the ID parameter as a function of the discrimination variable, x . The lowest cost at each x determines the ID. For example, in Figure 2a, at $x=0$, C_U as indicated by the dotted line, has minimum point cost. So $x=0$ is in the U region. As the ID parameter varies between relative proportions of 1:2:3 the identification varies Q:U:X. Figure 2a shows a case where an undecided region is derived for data suitable for treatment by classical discrimination, due to a lower cost for “no decision” than for an incorrect decision. Figure 2b shows how lowest cost results in more X being identified due to benefits (negative costs) for correct identification of X. Figure 2c shows how more Q are identified as a result of a high prior for Q. Figure 2d shows the counterbalancing effect of benefits for identifying X, and a high prior for Q. Detailed parameters are in Table 1.




Figure 3. Results for unequal variance, ($\text{sigq}=1.0$, $\text{sigx}=0.5$). Detailed parameters in Table 1, column 3. Plots of point costs (Equation 1) and the ID parameter as a function of the discrimination variable, x . The lowest point cost at each x determines the ID. For example, in Figure 3a, at $x=-1$, C_Q as indicated by the dashed line, has minimum point cost. So $x=-1$ is in the Q region. As the ID parameter varies between relative proportions of 1:2:3, the identification varies Q:U:X. Other than sigx , parameters are the same as for Figure 2. Note that due to $\text{sigq} > \text{sigx}$, Q is the correct ID for both small and large x . However, only a small total probability is associated with the large positive values of x . Note, also, the total absence of an X region for Figure 3c; C_X is nowhere the minimum; the ID level near $x=1.0$ is U.




Figure 4. Results for discrimination of Q, M, and X. Detailed parameters and numerical results in Tables 3 and 4. Plots of point costs (Equation 1) and the ID parameter as a function of the discrimination variable, x . The lowest point cost at each x determines the ID. From the lowest to the highest level of the ID parameter we have Q:U:M:X. For example, in Figure 4a, the lowest point cost for $x=0$ is C_M , so M is identified for $x=0$ where the ID parameter is at the third level. The principle result of this analysis is that the presence of M and Q complicates the discrimination of either from X. See the text for detailed discussion.




Figure 5. Results for discrimination of Q and X with a uniform prior distribution for X ; an analysis similar to outlier analysis. Detailed parameters and numerical results are given in Table 5. Plots of point costs (Equation 1) and the ID parameter as a function of the discrimination variable, x . The lowest point cost at each x determines the ID. From the lowest to the highest level of the ID parameter we have $Q:U:X$. For example, in Figure 5a, the lowest point cost for $x=0$ is C_U , so U is identified for $x=0$ where the ID parameter is at the second level. The principle result of this analysis is that true costs, if a uniform distribution is assumed for X , are higher than if one assumes the true distribution for X . See the text for detailed discussion.

OUTLINE OF EXPECTED COST OF CLASSIFICATION (ECC) THEORY

The action, k , which is performed in response to a measured discrimination parameter vector, x is determined by which action, k , has the lowest point cost, C_k , at x .

$$C_k(x) = \sum_i p_i \cdot f(x)_i \cdot c(k|i) \quad (1)$$

This results in minimum expected cost equal to the minimum C_k integrated over x .

In the formula, i represents true event type, *e.g.*:

explosion	X
earthquake	Q
mine blast	M

and k represents a possible action, *e.g.*:

decide	X
decide	Q
decide	M
decide	U (unidentified, no action)
decide	F (unidentified, fly satellite)

and

$c(k|i)$ are the costs for action k , given event type i , *e.g.*, $c(U|X)$

$f(x)_i$ is the probability distribution of x for event type i

p_i is the prior probability of event type i

To simplify to reach the classical linear discriminant one must assume:

- (1) k has same range, usually (1,2), as i , (*e.g.*, no U action)
- (2) $c(i|i)=0$ (*e.g.* no benefits, no different benefits for Q and X)
- (3) f_i are normal and have equal variance

Usually it is assumed that p_i are equal (but there are usually many more Q than X).

SUMMARY

The ECC method is a classical statistical procedure; calculations are simple and direct. The method minimizes cost instead of maximizing statistical power, and has many features characteristic of present procedures.

The method can consider:

- Multiple actions, including undecided
- Misclassification costs
- Correct-classification benefits
- System and political costs
- Multiple populations
- Unequal variances/empirical distributions
- Multiple dimensions
- Unequal priors: thresholds may vary with magnitude
- Uniform and truncated distributions if inadequate knowledge for, *e.g.*, X

Future work:

- Apply to actual data: *e.g.*, M_s : m_b and depth in 2D, including seismicity
- Develop realistic costs for various scenarios via interviews; apply
- Develop methods of applying sequentially, compare to multiple dimensions
- Develop software system to derive distributions from data and enter into ECC

DISTRIBUTION

California Institute of Technology
ATTN: Prof. Thomas Ahrens
Seismological Laboratory, 252-21
Pasadena CA 91125

Air Force Research Laboratory/VSBL
ATTN: Dr. Robert Raistrick
29 Randolph Rd.
Hanscom AFB MA 01731-3010

OATSD(NCB)/NT
ATTN: Mr. Patrick Wakefield
1515 Wilson Blvd., Suite 720
Arlington VA 22209

Pennsylvania State University
ATTN: Prof. S. Alexander & Prof. C. Langston
Department of Geosciences
537 Deike Building
University Park PA 16802

University of Colorado
ATTN: Prof. C. Archambeau, Prof. D. Harvey,
Dr. Anatoli Levshin, & Prof. M. Ritzwoller
Department of Physics/JSPC
3100 Marine Street
Boulder CO 80309-0583

Cornell University
ATTN: Prof. M. Barazangi
Institute for the Study of the Continents
Department of Geological Sciences
3126 SNEE Hall
Ithaca NY 14853

State University of New York, Binghamton
ATTN: Prof. F.T. Wu
Department of Geological Sciences
Binghamton NY 13901

ENSCO, Inc./APA Division
ATTN: Dr. D.R. Baumgardt & Dr. Z. Der
5400 Port Royal Road
Springfield VA 22151-2388

University of Arizona
ATTN: Dr. S. Beck & Prof. T.C. Wallace
Department of Geosciences
Gould Simpson Building
Tucson AZ 85721

Maxwell Technologies, Inc.
ATTN: Dr. T.J. Bennett & Mr. J. Murphy
Geophysics & Resource Technologies Group
11800 Sunrise Valley Drive, Suite 1212
Reston VA 22091

University of California, San Diego
ATTN: Dr. J. Berger, Dr. L. Burdick,
Dr. M. Hedlin, Prof. B. Minster,
Prof. J.A. Orcutt, & Dr. F.L. Vernon
Scripps Institution of Oceanography
Institute of Geophysics and Planetary Physics
9500 Gilman Dr.
La Jolla CA 92093

US Department of Energy
ATTN: Ms. Leslie A. Casey
NNSA/NA-22
1000 Independence Ave., SW
Washington DC 20585-0420

Multimax, Inc.
ATTN: Dr. W. Chan, Dr. I.N. Gupta,
& Mr. W. Rivers
1441 McCormick Drive
Largo MD 20774

Virginia Polytechnical Institute
ATTN: Dr. M. Chapman
Seismological Observatory
Department of Geological Sciences
4044 Derring Hall
Blacksburg VA 24061-0420

US Department of State
AC/VC
ATTN: Dr. M. Dreicer, Dr. C. Yeaw,
Dr. W. Leith, & Mr. R. Morrow
2201 C Street, N.W.
Washington DC 20520

University of Connecticut, Storrs
ATTN: Prof. V.F. Cormier
Department of Geology & Geophysics
U-45, Room 207
Storrs CT 06269-2045

DTRA/TDAS
ATTN: Dr. Anton Dainty, Dr. Jay Davis, &
Dr. R.A. Gustafson
8725 John J. Kingman Rd., MS6201
Ft. Belvoir VA 22060

DTRA/OSTN
ATTN: Dr. C. Gallaway & Dr. R. Goodwin
6801 Telegraph Rd.
Alexandria VA 222310

DTRA/CP
ATTN: Dr. Hopkins
6801 Telegraph Rd.
Alexandria VA 222310

Defense Technical Information Center
8725 John J. Kingman Road, Suite 0944
Ft. Belvoir VA 22060-6218

University of California, San Diego
ATTN: Dr. Catherine de Groot-Hedlin
Institute of Geophysics & Planetary Sciences
8604 La Jolla Shores Drive
San Diego CA 92093

Harvard University
ATTN: Prof. A. Dziewonski
Hoffman Laboratory
Department of Earth, Atmospheric & Planetary
Sciences
20 Oxford Street
Cambridge MA 02138

Boston College
ATTN: Prof. J. Ebel & Prof. A. Kafka
Department of Geology & Geophysics
Chestnut Hill MA 02167

Mission Research Corporation
ATTN: Dr. M. Fisk & Dr. W. Wortman
8560 Cinderbed Rd., Suite 700
Newington VA 22122

Southern Methodist University
ATTN: Dr. Henry Gray
Department of Statistical Science
P.O. Box 750302
Dallas TX 75275-0302

Pacific Northwest National Laboratories
ATTN: Dr. D.N. Hagedorn
P.O. Box 999, MS K5-12
Richland WA 99352

Lawrence Livermore National Laboratory
ATTN: Dr. J. Zucca
P.O. Box 808, L-205
Livermore CA 94551

MIT ERL E34-404
ATTN: Dr. David Harkrider
42 Carlton St.
Cambridge MA 02142-1324

New Mexico State University
ATTN: Prof. Thomas Hearn & Prof. James Ni
Department of Physics, 3D
Las Cruces NM 88003

California Institute of Technology
ATTN: Dr. Donald Helmberger
Division of Geological & Planetary Sciences
Seismological Laboratory
Pasadena CA 91125

Southern Methodist University
ATTN: Dr. E. Herrin & Dr. B. Stump
Department of Geological Sciences
P.O. Box 750395
Dallas TX 75275-0395

Sandia National Laboratories
ATTN: Mr. Pres Herrington
P.O. Box 5800, MS 0975
Albuquerque NM 87185-0975

St. Louis University
ATTN: Prof. R. Herrmann & Prof. B. Mitchell
Department of Earth & Atmospheric Sciences
3507 Laclede Ave.
St. Louis MO 63103

University of California, Berkeley
ATTN: Prof. L.R. Johnson & Prof. T. McEvilly
Earth Sciences Division
LBNL 90-2106, MS 90-1116
479 McCone Hall
Berkeley CA 94720

Massachusetts Institute of Technology
ATTN: Prof. T.H. Jordan
Department of Earth, Atmospheric, & Planetary
Sciences
77 Massachusetts Ave., 654-918
Cambridge MA 02139

Massachusetts Institute of Technology
ATTN: Dr. R. LaCross, M-200B
Lincoln Laboratory
P.O. Box 73
Lexington MA 02173-0073

University of Illinois, Urbana-Champaign
ATTN: Prof. F.K. Lamb & Prof. J. Sullivan
Department of Physics
1110 West Green Street
Urbana IL 61801

Oklahoma Geological Survey Observatory
ATTN: Dr. J. Lawson, Chief Geophysicist
Number One Observatory Lane
P.O. Box 8
Leonard OK 74043-0008

University of California, Santa Cruz
ATTN: Dr. T. Lay, Dr. S. Schwartz, & Dr. R. Wu
Institute of Tectonics, Earth Sciences Dept.
A232 Earth & Marine Sciences Building
Santa Cruz CA 95064

US Geological Survey
ATTN: Dr. John Filson
12201 Sunrise Valley Dr., MS-905
Reston VA 22092

Weston Geophysical Corp.
ATTN: Mr. James Lewkowicz
57 Bedford St., Ste 102
Lexington MA 02420

Georgia Institute of Technology
ATTN: Prof. L. Long
School of Geophysical Sciences
Atlanta GA 30332

Southern Methodist University
ATTN: Dr. G. McCartor
Department of Physics
P.O. Box 750175
Dallas TX 75275-0175

US Geological Survey
ATTN: Dr. A. McGarr
National Earthquake Center
345 Middlefield Rd., MS-977
Menlo Park CA 94025

Office of the Secretary of Defense
DDR&E
Washington DC 20330

Yale University
ATTN: Prof. J. Park
Department of Geology & Geophysics
P.O. Box 208109
New Haven CT 06520-8109

University of Cambridge
ATTN: Prof. Keith Priestly
Bullard Lab, Department of Earth Sciences
Madingley Rise, Madingley Road
Cambridge CB3 0EZ
UNITED KINGDOM

BBN Systems & Technologies
ATTN: Dr. J.J. Pulli
1300 N. 17th St., Suite 1200
Arlington VA 22209

DTR Associates/Weston Geophysical
ATTN: Dr. Delaine Reiter
325 West Main Street
Northborough MA 01532

Columbia University
ATTN: Prof. P. Richards & Dr. J. Xie
Lamont-Doherty Earth Observatory
61 Route 9W
Palisades NY 10964

Woodward-Clyde Federal Services
ATTN: Dr. C.K. Saikia
566 El Dorado Street, Suite 100
Pasadena CA 91101-2560

University of Southern California, Los Angeles
ATTN: Prof. C.G. Sammis
Center for Earth Sciences
University Park
Los Angeles CA 90089

Secretary of the Air Force
(SAFRD)
Washington DC 20330

University of California, Davis
ATTN: Dr. R. Shumway
Division of Statistics
1 Shields Ave.
Davis CA 95616-8671

AFOSR/NL
110 Duncan Avenue, Suite B115
Bolling AFB
Washington DC 20332-0001

Los Alamos National Laboratory
ATTN: Dr. H. Patton, Dr. S.R. Taylor,
Dr. C.L. Edwards, Dr. D.C. Pearson, &
Dr. Scott Phillips
P.O. Box 1663, MS D408
Los Alamos NM 87545

Massachusetts Institute of Technology
ATTN: Prof. M.N. Toksoz
Earth Resources Lab, 34-440
42 Carleton Street
Cambridge MA 02142

WINPACC/CA/FO
ATTN: Dr. L. Turnbull
New Headquarters Bldg., Room 4W03
Washington DC 20505

National Science Foundation
ATTN: Dr. Daniel Weill
Division of Earth Sciences, EAR-785
4201 Wilson Blvd., Room 785
Arlington VA 22230

NTNF/NORSAR
ATTN: Dr. Svein Mykkeltveit
Granaveien 33
P.O. Box 51
N-2007 Kjeller
NORWAY

Australian Geological Survey Organization
ATTN: Dr. David Jepsen
Cnr. Jerragomerra Ave. & Nindmarsh Dr.
Canberra, ACT 2609
AUSTRALIA

Atomic Weapons Establishment
ATTN: Dr. P. Marshall, Dr. David Bowers, &
Dr. Alan Douglas
Blacknest, Brimpton
Reading FG7 4RS
UNITED KINGDOM

University of Bergen
ATTN: Prof. Eystein Husebye
Institute for Solid Earth Physics, IFJ
Allegaten 41
N-5007, Bergen
NORWAY

AFTAC/CA(STINFO)/TT/TTD/TTE/TTR
1030 South Highway A1A
Patrick AFB FL 32925-3002