

**UNIFICATION OF STATISTICAL
AND ECONOMIC ANALYSIS**

Training Working Document No. 1

Prepared by



CIMMYT

TRAINING WORKING DOCUMENT

**UNIFICATION OF STATISTICAL
AND ECONOMIC ANALYSIS**

Training Working Document No. 1

Prepared by
Roger Mead
Consultant
in collaboration
with CIMMYT staff

CIMMYT
Lisboa 27
Apdo. Postal 6-641,
06600 México, D.F., Mexico

PREFACE

This is one of a new series of publications from CIMMYT entitled *Training Working Documents*. The purpose of these publications is to distribute, in a timely fashion, training-related materials developed by CIMMYT staff and colleagues. Some Training Working Documents will present new ideas that have not yet had the benefit of extensive testing in the field while others will present information in a form that the authors have tested and found useful for teaching. Training Working Documents are intended for distribution to participants in courses sponsored by CIMMYT and to other interested scientists, trainers, and students. Users of these documents are encouraged to provide feedback as to their usefulness and suggestions on how they might be improved. These documents may then be revised based on suggestions from readers and users and published in a more formal fashion.

CIMMYT is pleased to begin this new series of publications with a set of six documents developed by Professor Roger Mead of the Applied Statistics Department, University of Reading, United Kingdom, in cooperation with CIMMYT staff. The first five documents address various aspects of the use of statistics for on-farm research design and analysis, and the sixth addresses statistical analysis of intercropping experiments. The documents provide on-farm research practitioners with innovative information not yet available elsewhere. Thanks goes out to the following CIMMYT staff for providing valuable input into the development of this series: Mark Bell, Derek Byerlee, Jose Crossa, Gregory Edmeades, Carlos Gonzalez, Renee Lafitte, Robert Tripp, Jonathan Woolley.

Any comments on the content of the documents or suggestions as to how they might be improved should be sent to the following address:

CIMMYT Maize Training Coordinator
Apdo. Postal 6-641
06600 Mexico D.F., Mexico.

PRECISION IN NET BENEFIT ANALYSIS

This document is intended to be used in parallel with the Cimmyt Economics Training Manual "From Agronomic Data to Farmer Recommendations". References to tables or figures in that manual will be in the form Table M3.1

1. Basic Precision

All the agronomic data which are used in economic analysis are derived from experimental data. Consequently all such data is to some degree imprecise. That is, the values calculated as mean yields for particular treatments are estimates of the population mean yields for those treatments. The precision of the estimates is represented by the standard errors of the mean yields, which are usually obtained from an analysis of variance of the experimental data. For the data from the weed control experiment of Table M3.1 the analysis of variance is shown in Table 1.

Table 1: Analysis of variance of weed control data

Yields (kg/ha)			Analysis of Variance			
Location 1			Source	SS	df	MS
	Block 1	Block 2				
Treat 1	2180	2220	Blocks	84000	1	
Treat 2	3030	2570	Treatments	282300	3	
Treat 3	2440	2180	Error	66100	3	22033
Treat 4	3200	3060				
Location 2			Error Mean Square = 13200			
Location 3			Error Mean Square = 34900			
Location 4			Error Mean Square = 22900			
Location 5			Error Mean Square = 37000			

Since these error mean squares are fairly homogeneous we can calculate a combined analysis as follows;

Source	SS	df	MS
Blocks(for each location)	184350	5	36870
Treatments	2507920	3	835973
Locations	32045540	4	8011385
Treat x Location	553272	12	46106
Combined Error	174100	15	11607

We note the evidence of very large treatment and location effects and the non-negligible treatment by location interaction. However our only purpose now is to use the error mean square to estimate the precision of treatment means. The error standard deviation is 108 which is less than 5% of the overall mean. This is unusually low; in most on-farm experiments we might expect values between 10% and 20% and even higher values do occur.

The standard error of a difference between two means (derived from a total of 10 plots) is calculated as

$$\sqrt{(2(11607)/10)} = 48.$$

If we compare the treatment means using this standard error we have

Treatment	1	2	3	4	Standard error of difference based on 15 df
	1994	2444	2084	2600	48

Comparison of the treatment means shows that all treatment means differ significantly at the 5% level except for the difference between treatments 1 and 3 which is close to significance.

2. Precision of Net Benefits

All the calculations for Net Benefit Analysis and Marginal Analysis are based on the initial yield data and the precision of derived quantities may be calculated from the standard errors of the initial mean yields in the same way that the derived quantities are calculated. Thus in the calculations for the Partial Budget we first calculate the adjusted yields (80% of the initial mean yields for the weed control example); then the gross field benefits(x\$8/kg for the example); then the net benefits by subtracting the total costs that vary. The corresponding calculated standard errors for the example would be:

- | | |
|------------------------------|------------------------------------------------------------------------------------|
| 1) for adjusted yields; | 80% of that for initial yields |
| 2) for gross field benefits; | x\$8 of that for adjusted yields |
| 3) for net benefits; | unchanged since the total costs that vary are not subject to precision estimation. |

For the weed control example the results would appear as:

Treatment	1	2	3	4	Standard error(15df)
Mean yield (kg/ha)	1994	2444	2084	2600	48
Adjusted yield (kg/ha)	1595	1955	1667	2080	38
Gross field benefits (\$/ha)	12760	15640	13336	16640	307
Total variable costs (\$/ha)	2400	3875	3200	4675	0
Net benefits(\$/ha)	10360	11765	10136	11965	307

Notice that although the standard error in this example is small the significance of differences between treatments changes for the different estimated quantities. For the initial mean yields all differences are significant at 5% except for that between treatments 1 and 3. The same is true for the adjusted yields and the gross field benefits. However the net benefits show a different pattern, the differences between

treatments 1 and 3 and between treatments 2 and 4 both being small compared with the standard error. The net benefits for each of treatments 2 and 4 are still significantly different (5%) from those for each of treatments 1 and 3.

This change of significance for different quantities is very common and indeed we should expect it. The difference between treatments 2 and 4 for gross field benefits is 1000\$/ha. The same difference for net benefits is reduced by 800\$/ha because of the differential costs but the precision remains unchanged. Such a change in significance will appear surprising only if we allow ourselves to interpret "significance" as implying a real difference while "non-significance" is assumed to imply no difference. This is not a valid interpretation of significance which measures the strength of the evidence for a difference. In the example we should be strongly convinced that the gross field benefits for treatments 2 and 4 are different, but the size of the difference is not much more than the difference in the variable costs and so we have little grounds for believing that the net benefits for these two treatments are much, if at all, different.

3. Precision of Marginal Rates of Return.

The standard errors for marginal rates of return may also be derived from those for net benefits though now the standard errors for MRR comparing different treatments will be different. The MRR from treatment 1 to treatment 2 is calculated as the marginal net benefit divided by the marginal cost. We know the standard error of the marginal net benefit (the standard error of the difference between any two net benefit values) and so the standard error of the MRR is simply the standard error of the marginal net benefit divided by the marginal cost.

The MRR for treatment 1 to treatment 2 is calculated as

$$(11765 - 10360) / (3875 - 2400) = 1405/1475 = 0.95.$$

The standard error of this MRR is calculated as

$$307 / (3875 - 2400) = 307/1475 = 0.21.$$

The corresponding calculations for the MRR for treatment 2 to treatment 4 are:

$$\text{MRR} = (11965 - 11765) / (4675 - 3875) = 200/800 = 0.25$$

$$\text{Standard error} = 307 / (4675 - 3875) = 307/800 = 0.38.$$

We can also calculate the MRR for treatment 1 to treatment 4

$$\text{MRR} = (11965 - 10360) / (4675 - 2400) = 1605/2275 = 0.71$$

$$\text{Standard error} = 307 / (4675 - 2400) = 307/2275 = 0.13.$$

In significance terms we should be strongly convinced that the MRR for treatments 1 to 2 is greater than zero and moreover is also just about significantly (5%) greater than a critical MRR value of 0.5.

We can construct confidence limits for the MRR. For example the 95% confidence limits for the MRR for treatments 1 to 2 are

$$0.95 \pm 2.13 \times 0.21 \text{ giving } (0.50, 1.40).$$

The 90% confidence limits for the same MRR are

$$0.95 \pm 1.75 \times 0.21 \text{ giving } (0.58, 1.32).$$

The corresponding 95% limits for the other MRR values are:

MRR(2 to 4) $0.25 \times 2.13 \times 0.38$ giving (-0.56,1.06)

MRR(1 to 4) $0.71 \times 2.13 \times 0.13$ giving (0.43, 0.99).

Note that the MRR(1 to 4) is the most precise because it is based on the largest cost difference. However although it is even more convincingly different from zero than is the MRR(1 to 2) it is not so convincingly different from a critical MRR value of 0.5.

4. Graphical Representation of Precision

Much of the interpretation of precision information, and more generally of net benefits and marginal rates of return, is simplified by the use of graphical presentation of information. The basic structure of the graphical presentation is the net benefit/total variable costs diagram (as used in figure M4.1). For simplicity we first consider, in Figure 1, the case of two treatments, using treatments 1 and 4 of the weed control example. The precision information is presented as 95% confidence limits (other % confidence limits or simply standard errors could be used provided it is clear which form is being used).

The confidence limits for net benefits are shown in the form for marginal net benefits; that is for differences between the net benefits for two treatments. When we now consider the confidence limits for the MRR between treatments 1 and 4 we draw lines from the point value for the lower net benefit to the 95% confidence limits set about the upper net benefit. This correctly allows for the precision of the marginal net benefit and shows the slopes of the MRR corresponding to the confidence limits (0.43 and 0.99) calculated in the previous section.

More usefully we can display the MRR rates and their confidence limits for the sequence of intermediate (non-dominated) treatments, showing the MRR rates for each section of the net benefit curve. In Figure 2 the MRR rates for treatments 1 to 2 and for treatments 2 to 4 are shown, each with its corresponding 95% confidence limits. We can see how the range of credible MRR values is narrower when the cost difference is larger (compare also with Figure 1, where the cost difference is largest). The MRR(1to2) is sufficiently well estimated that the whole confidence interval is above the critical level of 0.5. In contrast, the MRR(2to4) is poorly estimated and includes a substantial range of negative MRR values.

5. Use of Critical MRR Lines

Another addition to the net benefit/variable costs diagram is the inclusion of lines representing the critical MRR value. In figure 3, which is based on data from a Nitrogen response experiment (Table M6.2) four treatments are compared for which the net benefits and variable costs are as follows:

		Net Benefits (\$/ha)	Total costs that vary (\$/ha)
Treatment 1	0kgN/ha	400	0
Treatment 2	40kgN/ha	486	30
Treatment 3	80kgN/ha	526	60
Treatment 4	120kgN/ha	535	85

Two lines representing the critical MRR value of 1.0, or 100%, are drawn. The first is drawn from the point representing the lowest cost/lowest yield treatment point and any treatment point above that line has an MRR value greater than 1.0 when compared with the lowest treatment. The second critical MRR line is

drawn backwards from the maximum yield treatment point and any treatment point in the triangle (A) is superior to the treatment giving the highest yield since the MRR from such a point to the maximum point must give an MRR rate less than the critical value of 1.0.

For this set of nitrogen response data there are two intermediate treatments (40kgN/ha and 80kgN/ha) either of which should be preferred to the maximum treatment. In order to assess, from the diagram, which of these should be chosen, we draw a third critical MRR line backwards from the higher (80kgN/ha) of the two alternatives and examine whether the lower treatment point falls above or below this new critical MRR line. It can be seen immediately that the lower point (40kgN/ha) is below the critical line drawn relative to the 80kgN/ha and hence we should choose the 80kgN/ha treatment as the best recommendation for farmers.

6. Confidence for Differences

Finally we can develop the ideas of the precision of estimated marginal net benefits and MRR rates to calculate the confidence that an MRR rate for any particular treatment is greater than the critical MRR rate. Consider the vertical gaps between the treatment points and the lower critical MRR line. Each gap represents the best estimate of the net benefit advantage over the minimum acceptable MRR. For the 40kgN/ha treatment the gap is

$$486 - (400 + 30) = 56.$$

Using the standard error of a difference of net benefits between two treatments we can calculate the confidence probability that the gap is genuinely greater than zero.

Suppose the standard error for a difference between the net benefit values for two treatments is 30 \$/ha (based on a large number of degrees of freedom since the mean benefits are derived from experiments at 20 locations). Then the ratio of the advantage over the critical MRR line divided by the standard of the marginal net benefit is:

$$\text{for 40 kg N/ha} \quad (486 - 430) / 30 = 1.87$$

$$\text{for 80 kg N/ha} \quad (526 - 460) / 30 = 2.20$$

$$\text{for 120 kg N/ha} \quad (535 - 485) / 30 = 1.67.$$

The confidence probabilities can be read from tables of tail probabilities for the Normal distribution (because the df are large); a short version of the appropriate table is attached as an Appendix, The results are:

Confidence for advantage with 40 kg N/ha being greater than the critical MRR rate is 0.9693 or almost 97%;

Confidence for advantage with 80 kg N/ha being greater than the critical MRR rate is 0.9861 or nearly 99%;

Confidence for advantage with 120 kg N/ha being greater than the critical MRR rate is 0.9525 or just over 95%.

We would usually be happy with any of these confidence levels but clearly the greatest confidence attaches to the advantage for 80 kg N/ha. Because the standard error for marginal net benefit is the same for all treatment pairs the confidence level depends only on the gap between the treatment net benefit and the corresponding point on the critical MRR line. Hence any treatment above the upper critical MRR line (in triangle (A)) will have a higher level of confidence than that for the maximum treatment. Essentially the treatment that would be preferred from traditional MRR arguments will also be preferred because it has the greatest confidence attached to its advantage over the critical MRR rate.

Finally we return to the weed control data to calculate the confidence for the advantages over critical MRR rates of 0.25, 0.5 or 0.75. The principles of the calculation are exactly as before except that we use tail probabilities based on the t-distribution because the degrees of freedom on which the standard errors are estimated are only 15. (Brief tail probabilities are tabulated in the Appendix).

	Treatment	
	2	4
Critical MRR rate = 0.25	$z = (11765-10729)/307$ $= 3.37$ Confidence P = 99.8	$z = (11965-10929)/307$ $= 3.37$ Confidence P = 99.8
Critical MRR rate = 0.5	$z = (11765-11098)/307$ $= 2.17$ Confidence P = 97.7%	$z = (11965-11498)/307$ $= 1.52$ Confidence P = 91.0%
Critical MRR rate = 0.75	$z = (11765-11466)/307$ $= 0.98$ Confidence P = 83.2%	$z = (11965-12066)/307$ $= -0.33$ Confidence P = 37.1%

For low critical rates both treatments show strong confidence (and if a zero critical MRR rate were considered treatment 4 would have a minutely stronger confidence). However for higher critical MRR rates the advantage of treatment becomes increasingly pronounced.

7. Precision and Variability

The methods in this paper relate to the use of information about the precision of estimation of yields, benefits, etc. We are not discussing the variability of results across locations, which is an important subject in its own right and is considered elsewhere. It may be noted, however, that the principles of the methods developed in this paper could be used with measures of variability in place of precision.

REPRESENTATION OF RISK

Assume that we are considering a recommendation to farmers to change from Treatment A to Treatment B on the basis of results from a substantial number of trials. The Marginal Rate of Return, calculated from the mean yields derived from the trial data, has been shown to provide an acceptable benefit from the proposed change. We would now like to examine the information about the variation of the returns over the set of trials. Essentially we wish to assess the probabilities of inadequate returns and, hopefully, identify those situations in which the inadequate returns occur.

The approaches described here will be illustrated for two sets of data. Initially we consider alternative methods for representing risk illustrating the methods for the data in Table 8.2 of the CIMMYT Economics Training Manual "From Agronomic Data to Farmer Recommendations". The second data set is from Zimbabwe and is drawn from training notes written by Allan Low. In this data set differential weights are allowed for the different observations.

1. The Variation in the Initial Data

Consider first the data for the two treatments, 0kgN and 80kgN at twenty locations. The Net Benefits are as shown.

Location	Net Benefits (\$/ha)		Location	Net Benefits (\$/ha)	
	0kgN	80kgN		0kgN	80kgN
1	441	655	11	542	562
2	511	647	12	512	681
3	383	277	13	285	291
4	391	610	14	387	578
5	250	593	15	375	230
6	322	619	16	494	661
7	490	660	17	485	660
8	458	600	18	295	480
9	180	162	19	485	683
10	250	612	20	463	260

To examine the joint pattern of pairs of benefits by location we plot a graph, shown in Figure 4, of the yield for 80kgN against the yield for 0kgN. The diagonal line through the origin represents equal benefits for the two treatments, so that points above the line are those for which the 80kgN treatment gives the higher benefits. This graph shows at once an unusual feature of this data, namely the division into two subsets: the main group of 15 observations towards the top of the graph and the smaller group of 5 towards the bottom and generally below the equality line. From this graph we observe also that the highest benefits are obtained from the 80kgN treatment and that the lowest benefits occur for both treatments with slightly more for 80kgN.

2. Risk Assessment Ignoring Location Pairing

Traditionally most of the methods of representing risk are based on the separate samples of data for the two treatments to be compared. The minimum returns analysis, described in the Training Manual, compares the averages of the 25% lowest benefits for each treatment. Two alternative representations utilising the complete data samples are the comparison of Cumulative Distributions and the Relative Risk Diagram.

2.1. Cumulative Distributions

The distribution of benefits for each treatment sample can be displayed in a Cumulative Distribution. This is constructed by first listing the sample values in increasing order. Using the sample for 0kgN we get

180,250,250,285,295,322,375,383,387,391,441,458,463,485,490,494,511,512,542.

Now we plot the proportion of values less than each value against that value. Thus there are no values less than 180. There is one value (out of 20) of 180 and the next value is 250 so that the proportion of values less than any value between 181 and 250 is 0.05. There are two values at 250 so the proportion less than any value between 251 and 285 is 0.15. From 286 to 295 the proportion of values less than that value is 0.20, and so on. The cumulative distributions for the 0kgN and 80kgN samples are shown in Figure 8.

By superimposing the cumulative distributions for the two treatments, as in Figure 8, we can compare the ranges and distribution patterns of benefits for the two treatments. The pattern at the lower level of benefits (\$150/ha to \$300/ha) is very similar for the two treatments. The 80kgN treatment shows a clear advantage of about \$150/ha at each proportion in the upper 70% of the distributions. This reflects the results for both the minimum returns analysis and average benefit increase in the Training Manual, showing little difference in benefits for the lowest 25% of results for each treatment but at the same time a large overall average benefit increase of \$126/ha.

2.2. The Relative Risk Diagram

A related technique displaying the advantage of one treatment against the other is the Relative Risk diagram (Mead et al; 1986). For this we require the two samples to be ordered together, but maintaining the identification of each value by its treatment.

0kgN	180		250	250			285		295	322	375	383	
80kgN	162	230			260	277		291					
0kgN	387	391	441	458	463		485	485	490	494	511	512	542
80kgN						480							
0kgN													
80kgN	562	578	593	600	610	612	619	647	655	660	660	661	681
0kgN													
80kgN	683												

The Relative Risk Diagram, shown in Figure 9, is constructed by moving through the joint ordering, counting along the horizontal axis of the graph the occurrences of the 0kgN sample, and up the vertical axis for the occurrences of the 80kgN sample. The levels of benefit are shown on the diagonal line (with a non-linear scale reflecting the benefit values which actually occurred in the sample). The diagonal line represents equal risks. For all critical levels of benefits in the range of interest this diagram shows the comparative probabilities, for each treatment, of benefits lower than any critical level

Thus, the first (lowest) benefit is 162, which is for the 80kgN treatment so that between benefit levels of 162 and 179 there is a 1/20 (=0.05) chance of a low yield for 80kgN and a 0/20 chance for 0kgN. After the first five values we have reached the benefit level of 250 and the chances are 2/20 (=0.10) for 80kgN and 3/20 (=0.15) for 0kgN. After ten values we have reached a benefit level of 295 and the chances are 0.25 for

both 0kgN and 80kgN. After twenty values we have reached a benefit level of 485 and the chances are 0.30 for 80kgN and 0.70 for 0kgN.

The main advantage of this diagram is that it emphasises the direct comparison of risks. The deviation of the relative risk curve from the diagonal shows the size of the difference in risk. Up to a benefit level of about \$300/ha the risks are similar for the two treatments, with slightly greater risks with 80kgN (the graph tending to be marginally above the diagonal). From benefits of \$300/ha up to about \$500/ha the risk of inadequate benefits increases steadily for 0kgN but hardly changes for 80kgN. At the critical benefit level of \$550/ha the relative risks are read off the graph by moving perpendicularly from the diagonal line and are 100% and 30%. The equivalent information can be extracted from Figure 8 making vertical comparisons between the two cumulative distributions, but that diagram does not display the magnitude of the risk differential so clearly.

Various patterns are possible in the relative risk diagram. A linear section of the relative risk curve through the origin, lying between one of the axes and the diagonal line represents a constant ratio of relative risk, the ratio of the risks being measured by the slope of the line. A section of the relative risk curve parallel to, but displaced from, the diagonal line represents a constant difference between the two risks. A horizontal, or vertical, section indicates rapid increase in one risk with increasing benefit level with no change in the other risk.

3. Risk Assessment Using Location Pairing

A disadvantage of the methods of the previous section is that they ignore the pairings of the results for each location. The methods could be applied equally well to data collected at two different sets of locations for the two different treatments. Not only do we fail to use the fact that the two treatments were both observed at each location but we formally assume that there is no relationship between the results for the two treatments. This is usually a most unrealistic assumption and, although the relationship shown in Figure 4 is, as has been noted previously, unusual, it does not suggest that the two sets of benefits should be assumed to provide independent information.

3.1. The relationship of Change to Current Practice

When the principal interest is in the change in benefit between two treatments, it is often beneficial to use an alternative form of Figure 4 in which the difference between the two benefits (80kgN - 0kgN) is plotted against the average benefit at each location. However from the point of view of the farmer deciding whether to change about from 0kgN to 80kgN the average of the two benefits is not a very appropriate measure. Instead we should plot the change in benefit against the net benefit for the 0kgN treatment at each location. This graph, which is a skewed rotation of Figure 4, is shown in Figure 7.

This graph displays how the advantages and risks of a change from 0kgN to 80kgN are distributed across the range of present practice, as represented by the benefits for each location of the 0kgN treatment. We can count directly from the graph the proportion of locations for which 80kgN gives a higher benefit than 0kgN, (16 out of 20), and the corresponding risk of 4/20 of failing to achieve an increase in benefit. If a critical MRR of 100% is assumed, as in the discussion of Table 8.2 in the Training Manual, then the minimum acceptable increase in net benefits would be \$60/ha and again we can count the proportion of locations giving increases of at least this level (14 out of 20). It may be helpful to draw the line of minimal acceptable increase on the graph, as shown in Figure 10.

We can also assess whether the risk/advantage is evenly spread over the range of benefits achieved from the 0kgN treatment. The six locations where negative changes or small positive changes occur are fairly evenly spread. The positive changes in benefit greater than the critical MRR do show a trend, however, with the largest advantages tending to occur at the lower 0kgN performance level.

Because Figure 7 contains all the points for the sample of individual locations the scatter of observations makes the pattern less easy to perceive. This is a situation where the smoothing techniques, developed in time series analysis of sequences of economic and other data, may be helpful. There is a wide range of smoothing methods available and the one used here is a relatively simple one chosen to illustrate the general philosophy rather because it provides the best results. A smoothed version of Figure 7 is constructed by considering a series of intervals of the horizontal axis and for each interval calculating the average change in that interval. Thus for the interval (300 - 400) there are five points

0kgN Benefit	Change in Benefit
322	+297
375	-145
383	-106
387	+191
391	+219

The average change in benefit over this interval is

$$(297 - 145 - 106 + 191 + 219)/5 = 91.2.$$

By calculating the average change for intervals of width 100 centred on 175, 200, 225, 250, 275, 300, 325, 350, 375, 400, we produce, in Figure 8, the desired smoothed version of Figure 7.

The result shows a pronounced peak (caused mainly by two locations) at a 0kgN benefit level about \$250/ha followed by quite a strong dip around 0kgN benefit levels around \$350 to \$400. This, in turn, leads to a more steady plateau of \$100 change in benefit for the higher levels of 0kgN benefit. At both ends the information is, of course, based on very few values.

There are three main influences on this form of diagram. The fact that we are plotting (y-x) against (x) will tend to produce a negative slope arising from the negative correlation of random errors in the absence of systematic patterns. At very low yield environments neither treatment can produce good yields and the yield difference must be small, with the change in benefit being negative. Similarly it might be expected that at very high yield levels the change in benefit would tend to decrease. In the middle range of the curve a horizontal section indicates that the improved treatment is maintaining its advantage over the standard, in defiance of the natural tendency to negative slope. If the slope tends to be positive then the improved treatment shows an even greater advantage against the natural negative correlation.

Thus the surprising aspect of Figure 8 is not the initial slightly negative value when both treatments produce poor yields nor the early peak, but the dip between the peak and the final plateau. The three locations which cause this are seen more clearly in Figure 7 and represent a proportion of medium to good environments where the improved treatment simply does much worse than elsewhere. There is no trend but a clear split into good and bad results for the improved treatment.

3.2 Cumulative Expressions of Advantage and Risk

It is sometimes more meaningful to discuss the average performance for all locations below a particular threshold than to consider each location separately. The idea of the cumulative distribution in section 2.1 exemplifies this concept (in contrast to a histogram of benefits). We therefore consider another modification to Figure 7 to consider the cumulative average benefit increase instead of the individual values.

To construct the cumulative mean advantage graph, shown in Figure 9, we order the locations in increasing value of 0kgN benefit. For each location we then calculate the average change in benefit for all those locations with the same or lower value of benefits from 0kgN. The resulting cumulative average advantages show how the advantage of the improved treatment develops as the yield environment is gradually improved. The first few steps of the calculation are shown.

Locations in order of	0kgN benefit	Changes to be averaged	Mean
9	180	-18	-18
5 & 10	250	-18, +362, +343	+229
13	285	-18, +362, +343, +6	+173
18	295	-18, +362, +343, +6, +185	+176

After the initial small negative change there is a sharp rise to a value over 200 and then the cumulative mean advantage settles down, quite quickly to its final level around 130.

Finally, as well as calculating the cumulative mean advantage in Figure 9, we can examine the cumulative risk, calculated in a parallel manner. As before the locations are listed in order of ascending benefit at 0kgN. For each level of 0kgN net benefit we count the proportion of those locations with the same or lower value of benefits from 0kgN which give a positive advantage. A second expression of this form of risk is to count the proportion of such locations giving a change greater than the 60\$/ha required to surpass the critical MRR. The calculations are again shown for the first steps of the graph.

Locations in order of 0kg N benefit	0kg N benefit	Changes to be counted	Proportions	
			above zero	above 60
9	180	-18	0/1	0/1
5 & 10	250	-18, +362, +343	2/3	2/3
13	285	-18, +362, +343, +6	3/4	2/4
18	295	-18, +362, +343, +6, +185	4/5	3/5

The resulting graph, with both cumulative proportions is shown in Figure 10.

Both cumulative proportion curves, after the settling-down process of the first two points, show very steady proportions of about 0.7 and 0.6 respectively, slowly tending upwards to their final values of 0.8 and 0.7. Thus, like the cumulative mean advantage graph the pattern is of a consistent average level after very few initial variations and with a slight late upward trend.

In both Figures 9 and 10 we would be pleased to find that the cumulative values settled down quickly. The cumulative advantage curve in Figure 9 is subject to the same tendency to a downward trend caused by the negative correlation between $(y-x)$ and (x) as was expected in Figure 8. The interpretation of patterns in Figure 9 are, generally, very much the same as in Figure 8 except that the cumulative plotting should reduce the scope for meandering. In fact we do get the negative trend in Figure 9 suggesting that there is a generally consistent level of advantage, the more extreme peak in Figure 8 being caused by the exceptionally high advantage of the improved treatment in those locations where the 0kgN benefit was low.

The cumulative proportion graph in Figure 10 should also be expected to settle to a steady level and there should not be the same expectation of a negative trend because the correlation does not apply to the same extent for counts. Unlike the cumulative advantage curve, the cumulative proportion curve must show, at least, some jerkiness because each successive observation must go up or down from the previous one.

4. The Zimbabwe Data (Using Weighting)

The second example set of data is from a worked example attached to teaching notes on "Application of Risk Decision Theory to Net Benefit and MRR Analysis" by Allan Low. There are eighteen sample locations regarded as representative of a potential domain population. Within the total population it is estimated that the proportions of "Good", "Medium" and "Poor" locations are 30%, 50% and 20% respectively. The eighteen sample locations are classified as being six from the "Good" locations, six from the "Medium" locations and six from the "Poor" locations. It is therefore decided to weight the observations in the three groups, according to these estimated overall proportions.

Thus the six good sample locations should predict for 30% of the total population; the six medium sample locations for 50% of the total population; the six poor sample locations for 20% of the total population. The practical effect of this differential importance is achieved by allocating weights of 3, 5 and 2 to the observations from the three groups to indicate their relative importance.

This use of weighting is unusual but, provided there are clear grounds for assessing the weights, it offers a method of using knowledge about the degree of typicalness of sample values. Most commonly we have a sample which is "randomly" chosen, if randomness is interpreted in the context of various practical restrictions. We then have to treat each sample as equally informative. However we may know that our sampling proportion has been different in different areas, deliberately or by the accidents of loss of results. Weighting is an attempt to correct unbalanced sample proportions. Not weighting would imply acceptance of the sample proportions as representative of the population proportions. Weighting may also be appropriate as a mechanism for using subjective judgements about the extremeness or unusualness of different years.

The data for benefits and benefit changes in the 18 locations are shown.

Class	Net Benefit		Change in Benefit	Weight
	0kgN	60kgN		
Good	219	432	+213	3
	359	460	+101	3
	84	243	+159	3
	275	498	+223	3
	407	508	+101	3
	161	479	+318	3
Medium	269	251	-18	5
	221	262	+41	5
	434	640	+206	5
	246	339	+93	5
	172	240	+68	5
	255	324	+69	5
Poor	14	-75	-89	2
	101	164	+63	2
	29	9	-20	2
	74	102	+28	2
	191	154	-37	2
	184	94	-90	2

The same seven forms of graphs (numbered 11, 12, 13, 14, 15, 16 and 17) are presented for this second set of data and additional explanation will be given only where the weighting requires modification of the construction procedure. Discursive comments will also be restricted to points of contrast with the previous example.

Figure 11: Joint Variation of the Two Treatments.

The weights for each location point are represented by numbers (3,5,2). The pattern is much more homogeneous than in the previous example with points around the diagonal of equal benefit for low benefit levels and tending increasingly above the diagonal for higher benefit levels.

Figure 12: Cumulative Distributions

The counting process includes the weights so that the plot shows the proportion of the total weight (60) associated with the locations giving benefits less than each value. The first few calculations are shown.

Ordered values for 0kgN	14	29	74	84	101	161	172
Weights	2	2	2	3	2	3	5
Cumulative Proportion of Weight	2/60	4/60	6/60	9/60	11/60	14/60	19/60

Figure 13: Relative Risk

Modified in exactly the same way as Figure 12. The result, though, is more clearly different with a fairly consistent differential of about 0.25 between the risks for 60kgN (lower risk) and 0kgN (higher risk) once the level of benefits is over 200.

Figure 14: Change v. Current Practice

Modified as for Figure 11. Comments as for Figure 11.

Figure 15: Smoothed Changes

Instead of using simple averages in each interval we use weighted averages. The resulting graph shows less variation of trend than was achieved for the previous example, though the scatter in Figure 10 has not been fully smoothed out. There is a clear initial negative benefit change, which is converted into a positive benefit change by the stage of a current net benefit level of 100; the change then maintains its level with a very slight upward trend. As mentioned earlier the negative change for the lowest environmental conditions is quite common. There is no evidence here of the negative correlation, the change in benefit showing a steady slight increase.

Figure 16: Cumulative Mean Advantage

As for Figure 15 weighted averages are used instead of simple averages. Also a much simpler pattern showing an initial loss at low 0kgN benefit values changing steadily to a gain of about 75 maintained at that value.

Figure 17: Cumulative Proportion

Modification exactly as for Figure 12. Again a smoother and simpler pattern of change.

DISCRETE AND CONTINUOUS ANALYSIS

Introduction

One important group of experimental research investigations, with consequent recommendations, is concerned with the effects of differing levels of a quantitative input factor. The most common example is the investigation of the effect of nitrogen fertilisers and the discussion here will be written in the context of a nitrogen response relationship. An example of a detailed analysis of nitrogen responses is contained in document 2B (Mead, 1990c).

A major decision for the experimenter and analyst is whether to develop the research in terms of a discrete or continuous model. Obviously the underlying model is a continuous one. Any amount of nitrogen could be applied. But in an experiment or for a recommendation a finite set of alternatives will inevitably be considered. I think there are three separate stages of experimentation and analysis, for each of which we have to ask the question "Discrete or Continuous?",

- (i) the choice of experimental levels of nitrogen,
- (ii) the analysis of the experimental data,
- (iii) the calculation of net benefits, marginal rates of return, and the framing of recommendations.

I believe that these questions involve different principles and, although linked, should be considered independently. Thus it is not, for example, necessary that the sets of alternative levels for the choice of experimental treatments and for possible recommendations should be identical.

1. The Choice of Experimental Levels

Obviously the experimental levels must be a discrete set. Statistical theory is quite clear that to investigate the response to a quantitative factor the maximum information is achieved by using the minimum number of levels compatible with the requirement of being able to estimate the appropriate response function. In practice most response functions include three parameters and to allow estimation three levels of N are needed; to also be able to assess whether the response model provides an acceptable fit we need a fourth level. More levels dissipates the information and gives less efficient estimation of the response function.

The second choice is which levels and again statistical theory indicates clearly that the levels should be chosen to cover as wide a range as possible, subject to the proviso that the form of the response curve should be credible over the entire range of levels. In practice this usually means that we take four equally-spaced levels starting with zero with the third level about the expected (economic) optimum N level.

2. The Analysis and Summary of the Experimental Data

Since the response being investigated is truly continuous and the experimental levels are simply a representative sample the summary should be in terms of a fitted response function. Statistical analysis through comparison of pairs of treatments with assessment of significance of treatment differences makes no logical sense since if there is any pattern of response at all no Null Hypothesis of a zero difference is credible.

The most suitable response curves are those which allow for non-symmetry of the response. If a quadratic response is preferred then a quadratic in the square-root of N will probably provide a better fit (measured in terms of the residual mean square or $(1-R^2)$, rather than R^2) than the ordinary quadratic. Inverse polynomials have been found to give generally better fits than ordinary polynomials. From the fitted

response, with estimates of precision, we can predict the response at any level, including any of the experimental levels, more accurately than from the individual experimental treatment mean yields (again see document 2B, section 2).

3. Calculation of Net Benefits, etc.

The assessment of possible recommendations in terms of a set of discrete alternatives has obvious intuitive appeal. It can be presented as a sequence of decisions each of which involves a sufficient change for the precision of the marginal rate of return to be reasonably precisely estimated (see document 1A). However the set of possible alternatives is not necessarily the same as the set of experimental levels. We might well wish to consider more alternatives for recommendation than we should wish to use as experimental levels.

The calculation of net benefits will be considerably more accurate if the predicted values from the response curve fitted to the experimental mean yields are used. This is essentially a smoothing argument. For example, it would be generally accepted that the yield response to a sequence of equally-spaced levels of N will follow a pattern of decreasing increments. Experimental treatment means, because of their associated standard errors will frequently produce successive yield increments which appear not to adhere to this expectation. The standard error of the fitted estimate of the yield difference between the two middle levels of four equally-spaced levels is less than half that of the observed difference between the experimental treatment mean yields.

4. Summary

We should use

- (i) three or four widely spread levels of N for the experimental treatments,
- (ii) an appropriate response curve to summarise the experimental data,
- (iii) the predictions from the fitted response curve when assessing benefits, and
- (iv) a set of discrete alternative levels of N when determining the best recommendation.

Figure 1




Figure 2





Figure 3



80 kg N
Net benefit

Figure 4

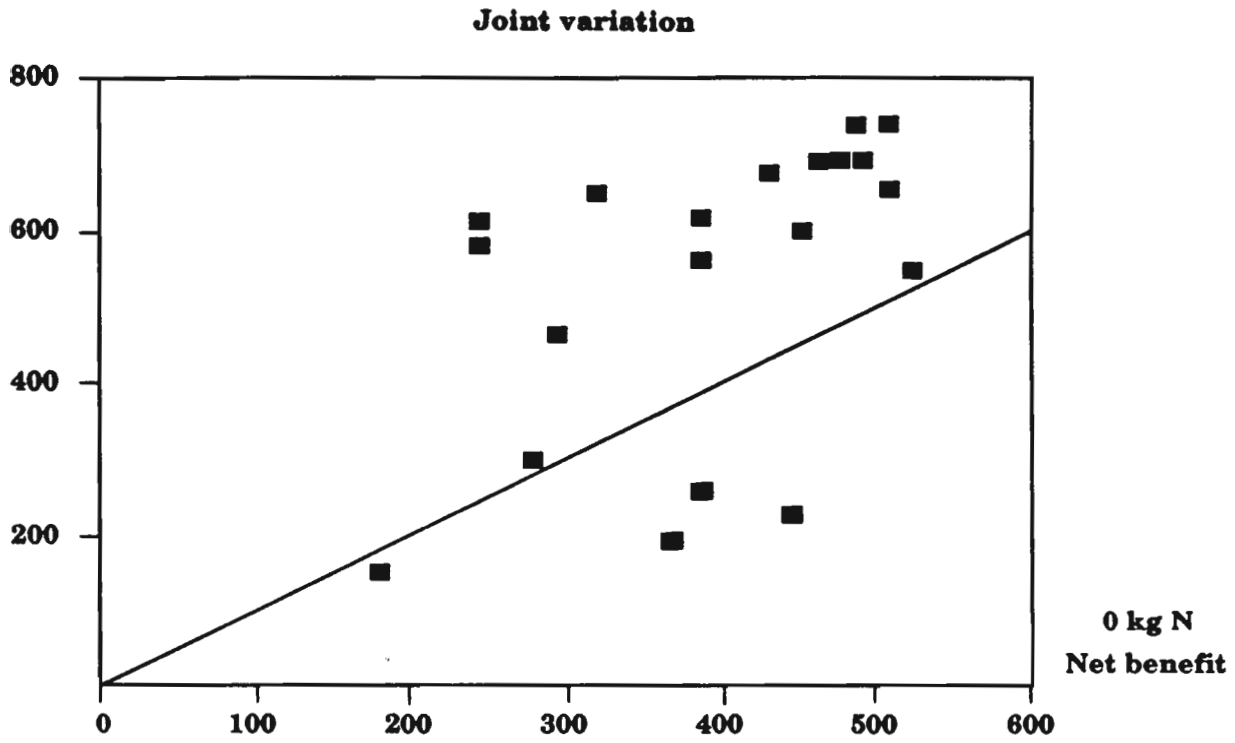


Figure 5

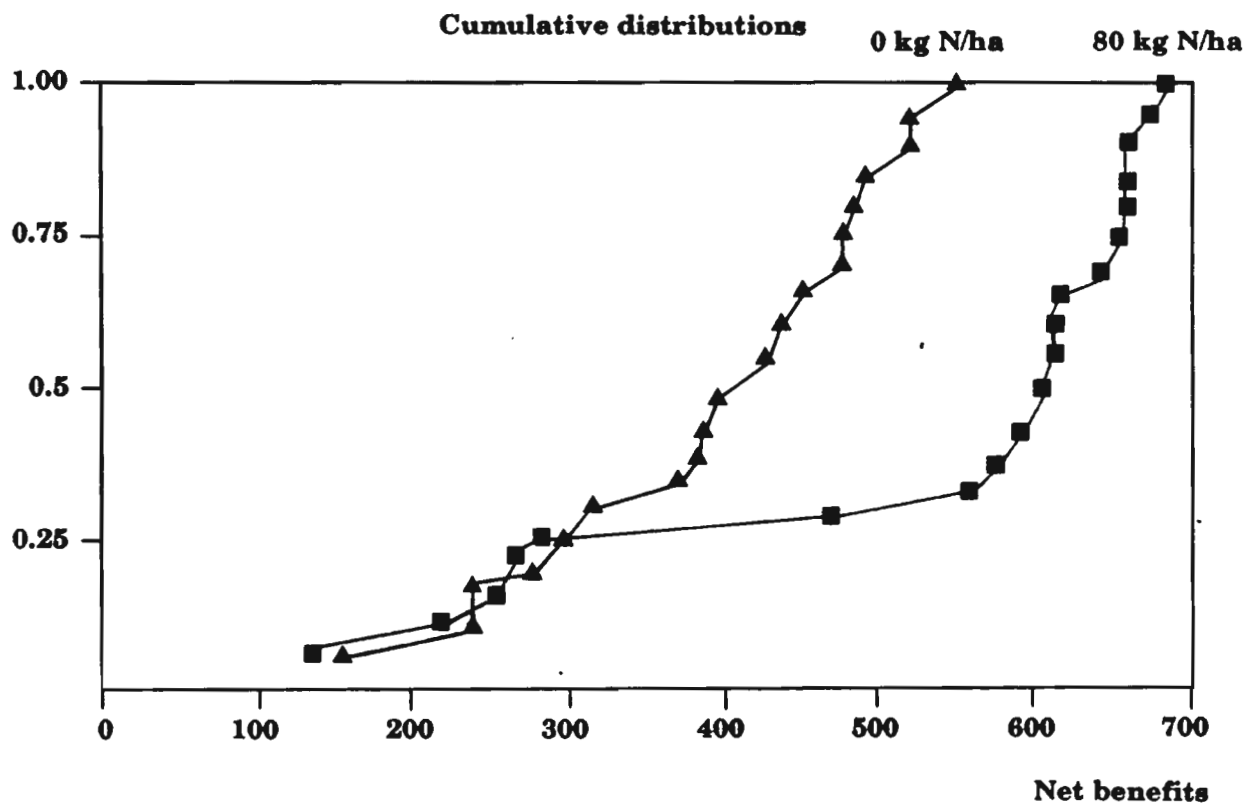


Figure 6

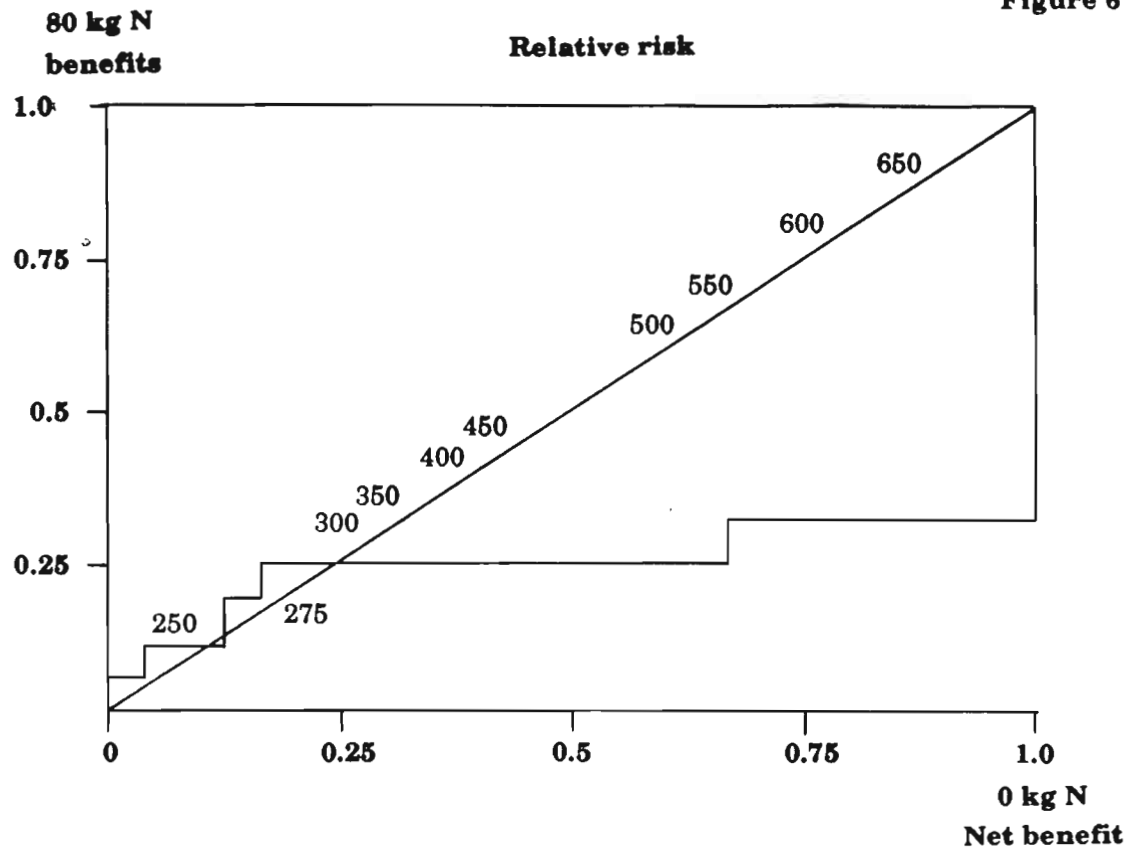


Figure 7




Figure 8

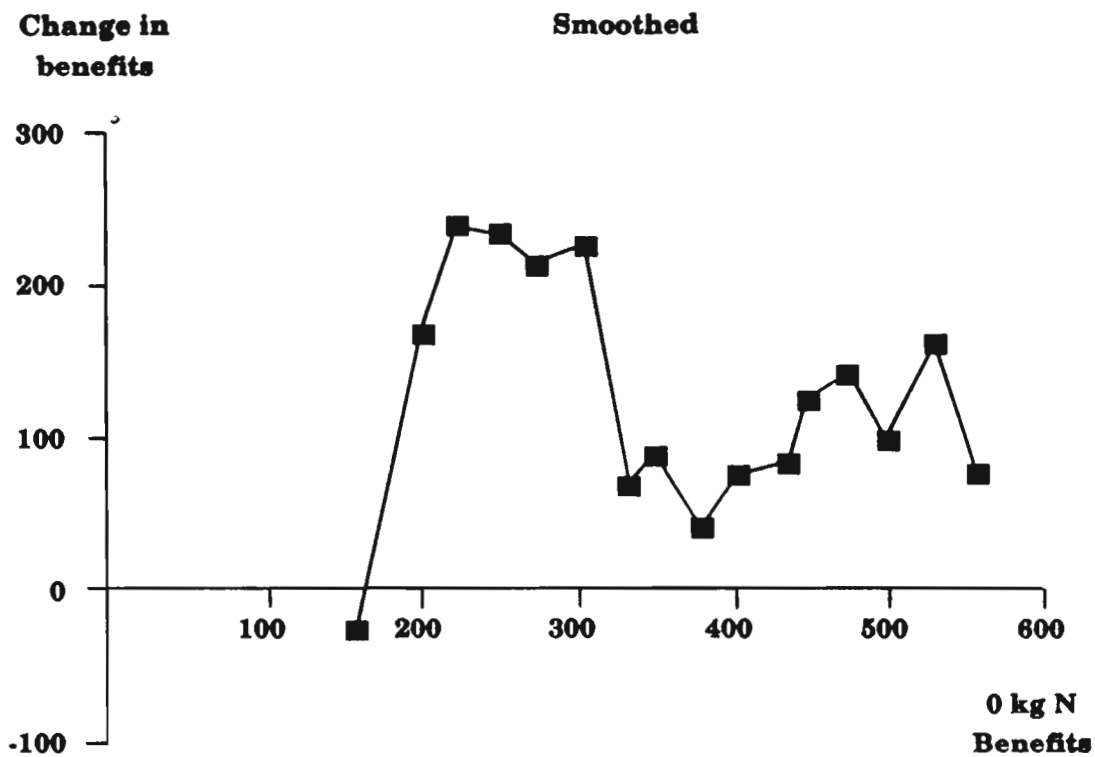
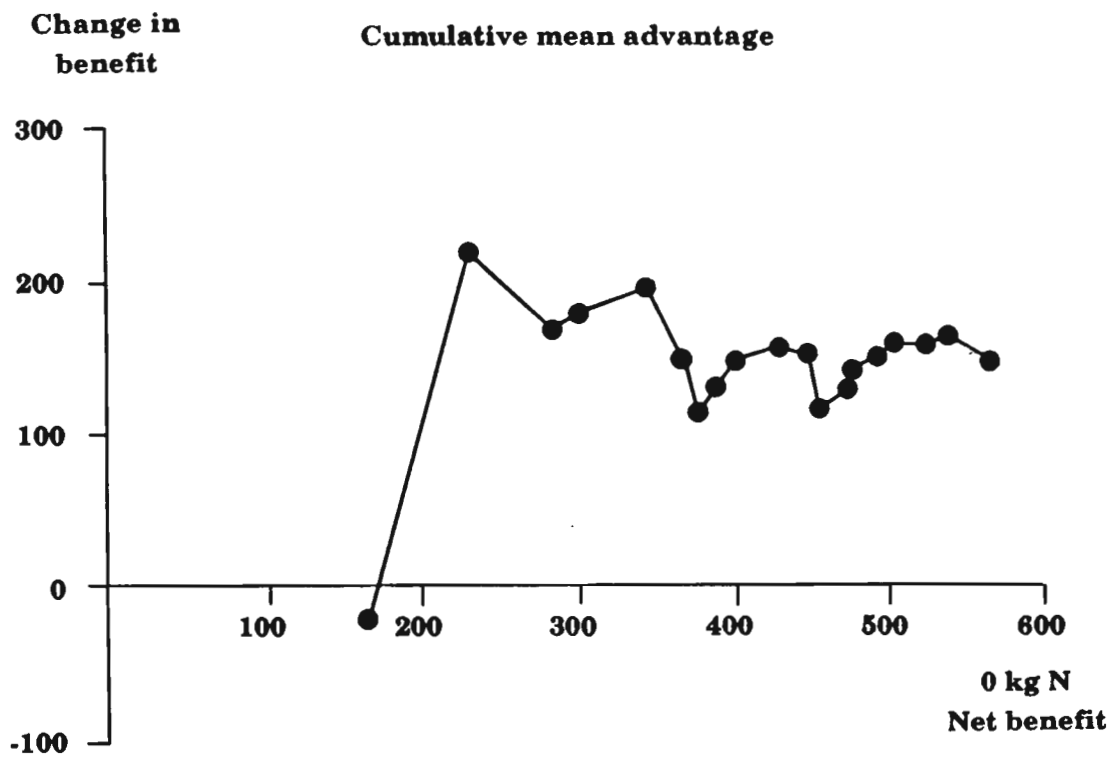



Figure 9



Cumulative proportion < value

Figure 10



60 kg N
Net benefits

Figure 11




Figure 12

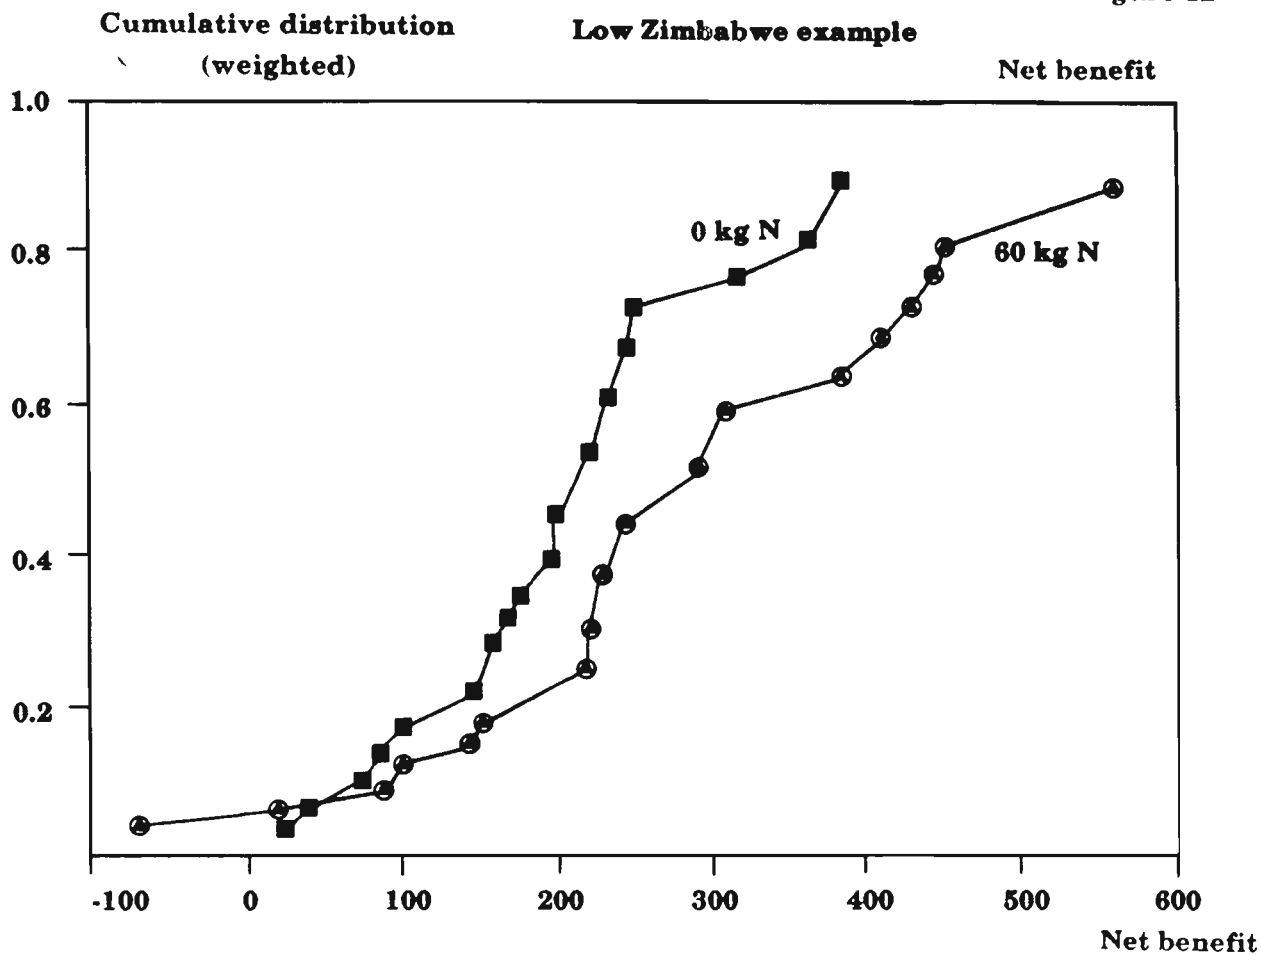
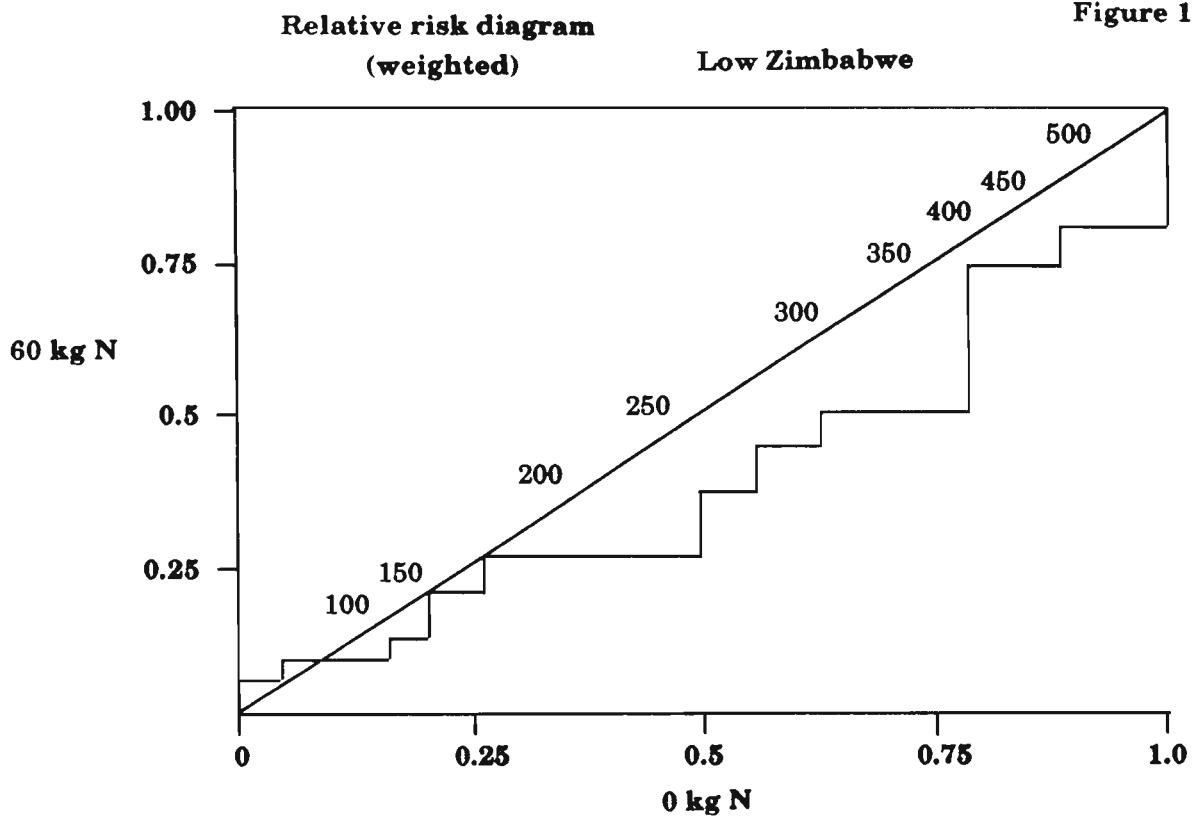


Figure 13



Change against 0 kg N
(weighted)

Figure 14




Figure 15

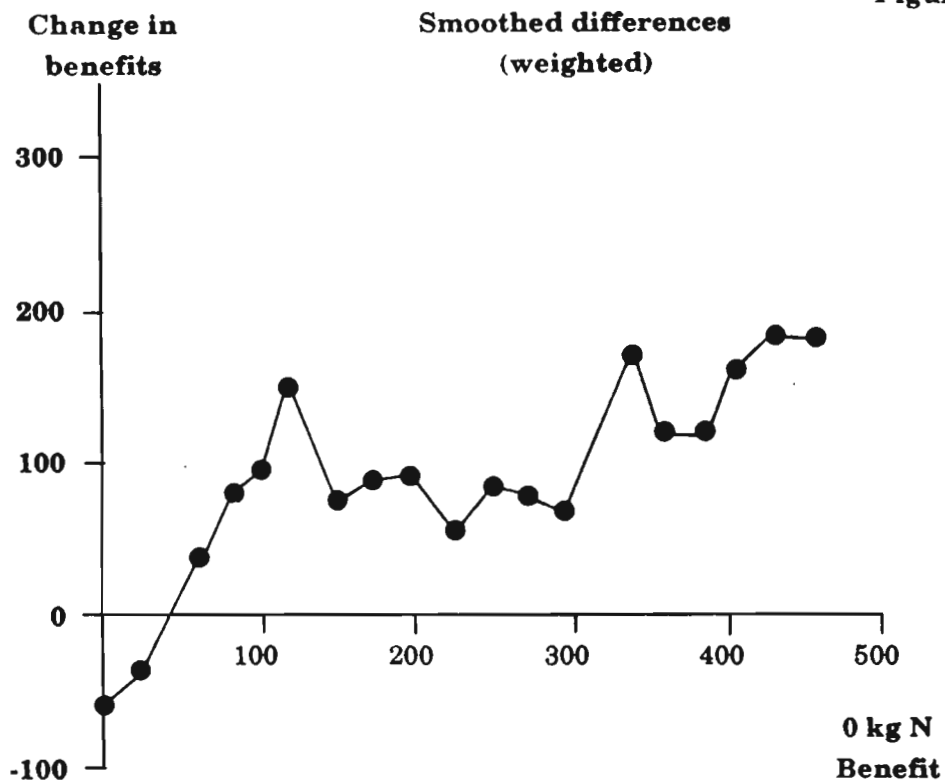


Figure 16

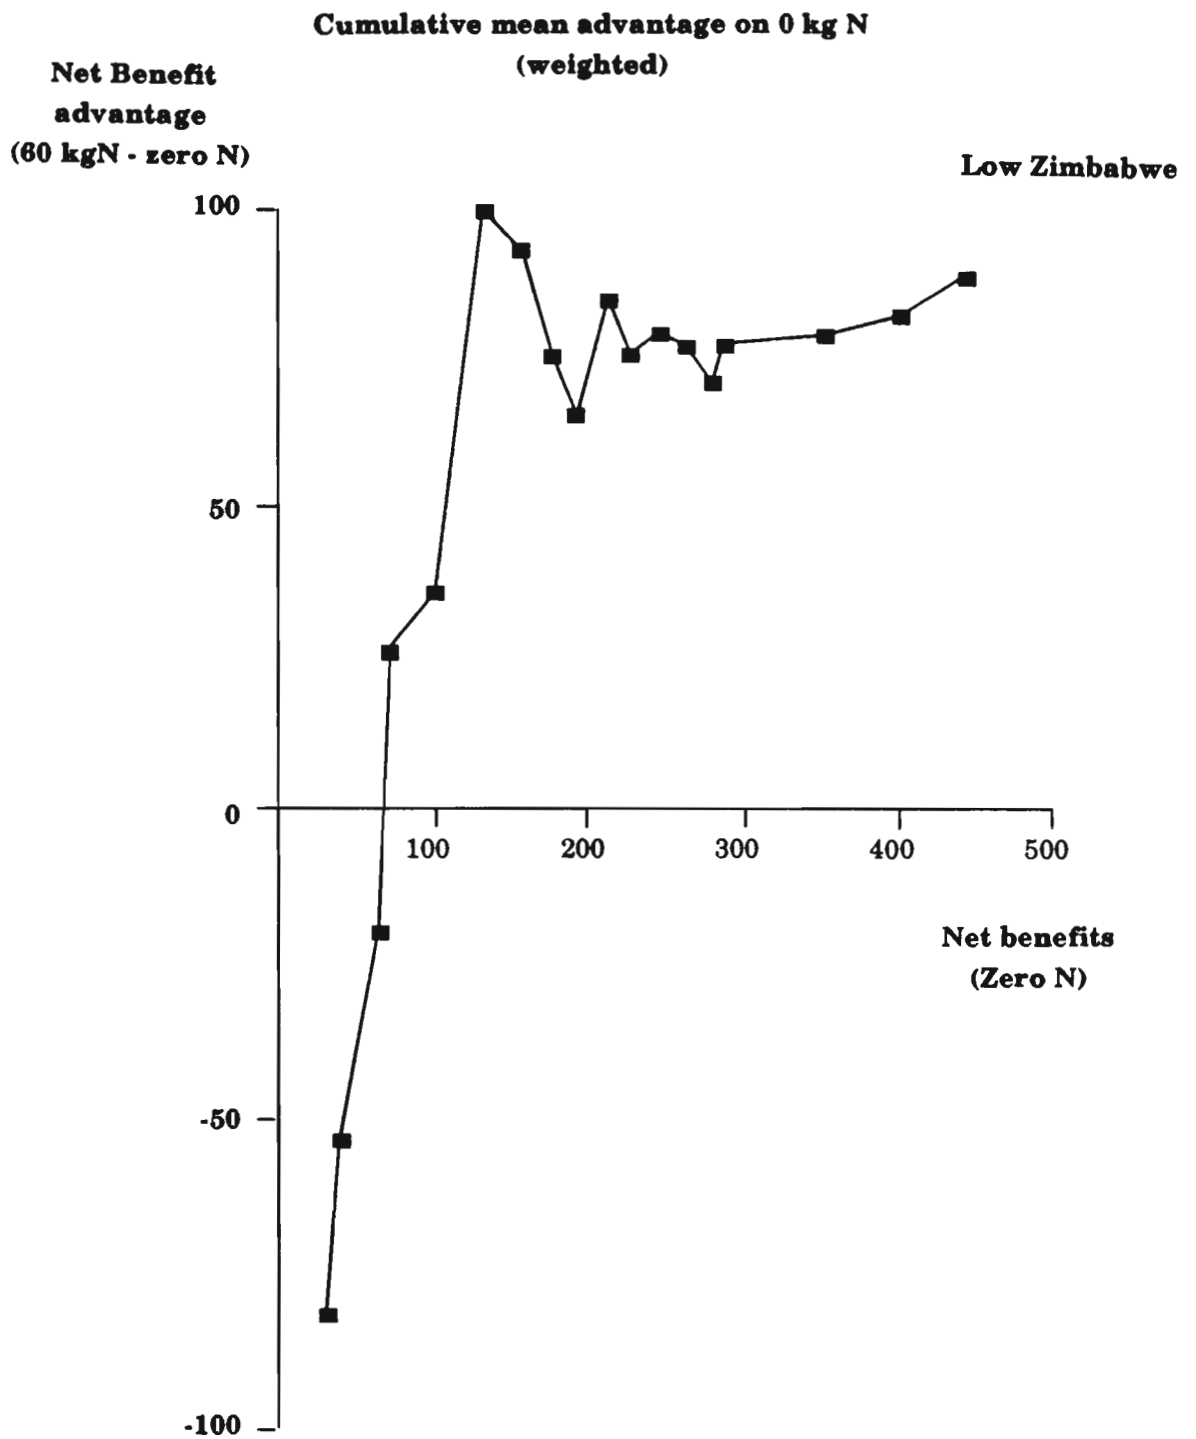
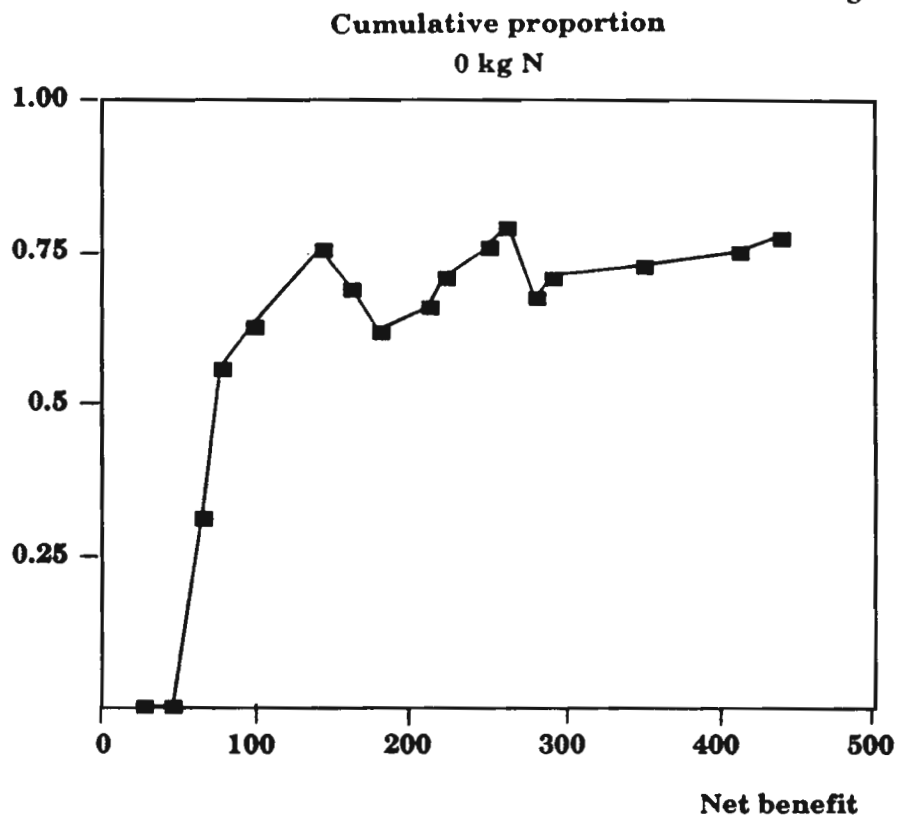


Figure 17





Centro Internacional de Mejoramiento de Maíz y Trigo
International Maize and Wheat Improvement Center
Lisboa 27, Apartado Postal 6-641, 06600 México, D.F. México