THE USE OF HALF-NORMAL PLOTS IN THE INTERPRETATION OF TWO-LEVEL FACTORIAL EXPERIMENTS

By

VADIS GLENN GODBEY

Bachelor of Arts

Berea College

Berea, Kentucky

1951

Submitted to the Faculty of the Graduate School of
the Oklahoma State University
in partial fulfillment of the requirements
for the degree of
MASTER OF SCIENCE
May, 1963

THE USE OF HALF-NORMAL PLOTS IN THE INTERPRETATION

OF TWO-LEVEL FACTORIAL EXPERIMENTS

APPROVED:

David L. Hicks

Report Adviser

Herbert Schol

Robert MacVicar

Dean of the Graduate School

PREFACE

The inspection and criticism of the data from 2^p factorial experiments are always necessary. A rule of inference is given for detecting a small number of real effects or interactions in the presence of a majority of error contrasts. This rule uses the ratio of the largest order-statistic of the contrasts to one whose expected value is nearest the standard error of contrasts. A sensitivity function for this ratio is developed. The possibility of discovering the presence of bad values in the raw data is also studied.

I wish to express my sincere gratitude to Dr. David L. Weeks for his suggestion of the topic of this report and for his interesting comments on the subject. Appreciation should also be expressed to Professor Herbert Scholz, Jr., for serving on my advisory committee. Certainly, I am indebted to Mrs. Mary Jane Walters for typing this report in her off duty hours.

V. G. G.

TABLE OF CONTENTS

LIST OF FIGURES

CHAPTER I

INTRODUCTION

The idea of plotting the empirical distribution of the usual set of orthogonal contrasts from a 2^p factorial experiment on a special grid as an aid in the criticism and interpretation of the experiment seems to have been that of Cuthbert Daniel. Even if the idea were not originally his, he has adopted it to the point that he is probably the best known exponent of the use of half-normal plots in the interpretation of suitable data. Most of the published literature on the subject of half-normal plots is by Cuthbert Daniel and his friend, Allan Birnbaum. Daniel[5] reports that J. W. Tukey has developed a half-normal grid in which the half-normal line is a horizontal straight line and the logarithms of the absolute values of the contrasts are used as one coordinate. It is believed that a paper on Tukey's half-normal grid has never been published since it could not be found in the available literature and since Daniel indicated that his information on this half-normal grid was received from Tukey in personal communications. This paper, then, is devoted primarily to a discussion of the half-normal plot as developed and used by Cuthbert Daniel and Allan Birnbaum. More specifically, this paper discusses the use of the half-normal plots in the interpretation of two-level factorial experiments.

CHAPTER II

CONSTRUCTION OF A HALF-NORMAL PLOT

In a 2^p factorial experiment there are $2^p - 1$ contrasts. Each of these $2^p - 1$ contrasts is orthogonal and is composed of 2^{p-1} observations minus 2^{p-1} observations. If the original observations have errors which are independent and normally distributed with mean zero and variance σ^2, then the $2^p - 1$ contrasts are independent and normally distributed with variance $2^p \sigma^2$. If the experiment is a null experiment, that is, there are no real effects or interactions, then the expected value of the contrasts is zero for each contrast and the contrasts are independent and normally distributed with mean zero and variance $2^p \sigma^2$. However, even though the expected value of each contrast is zero, we know that since the original observations have errors, the $2^p - 1$ contrasts will not be, in most cases, equal to zero. In fact, they may be either positive or negative and their sign is arbitrary. This indicates that maybe we should look at the absolute values of the contrasts. If there were no true effects or interactions, these contrasts will follow the positive half-normal distribution whose density is

$$f(x) = \sqrt{\frac{2}{\pi\sigma_c^2}} \exp\left(\frac{-x^2}{2\sigma_c^2}\right)$$

where $0 \leq x < \infty$ and $\sigma_c^2 = 2^p \sigma^2$ is the variance of a contrast.

This suggests that we can compare the $2^p - 1$ contrasts with this distribution. Of course, this comparison may be done in several ways. One of these ways is to graph the cumulative sample distribution of the contrasts against the cumulative half-normal distribution. It is this method, which seems to be a very informative one, that is the subject of this report.

Normal probability paper is so designed that the cumulative distribution of a normally distributed chance variable appears as a straight line. This is achieved by means of a non-linear transformation of the vertical scale of the graph of the cumulative normal distribution curve. The result is a paper in which the h (horizontal) scale is arithmetic and the v (vertical) scale, which ranges from $-\infty$ to $+\infty$, is labeled so that v equals zero is marked 50%, v equals 1.282 is labeled 90%, v equals 1.645 is labeled 95%, etc. The variable v is not normally printed on the paper, but only a sequence of the values of P where

$$P = \text{label for ordinate } v = \int_{-\infty}^{v} \frac{1}{\sqrt{2\pi}} \exp\left(\frac{-x^2}{2}\right) dx, \quad -\infty < x < +\infty.$$

Let us now consider a random sample $x_1, x_2, x_3, \cdots, x_n$ from a normal density function with mean zero and variance σ_c^2. Let k be a number such that P of the x's are less than or equal to k. Then

$$P = \int_{-\infty}^{v} \frac{1}{\sqrt{2\pi}} \exp\left(\frac{-x^2}{2}\right) dx \doteq \int_{-\infty}^{k} \frac{1}{\sqrt{2\pi\sigma_c^2}} \exp\left(\frac{-x^2}{2\sigma_c^2}\right) dx$$

$$= \int_{-\infty}^{k/\sigma_c} \frac{1}{\sqrt{2\pi}} \exp\left(\frac{-u^2}{2}\right) du; \quad -\infty < u < +\infty$$

where $\doteq$ means "is estimated by". Hence

$$v \doteq \frac{1}{\sigma_c} k \quad .$$

To put this in words, the plot of v against k for different P values will be linear through the origin with slope $1/\sigma_c$. Also, the value of k that corresponds to a v value of unity is equal to σ_c.

In the present case, we are concerned with the half-normal distribution. We will, therefore, ignore the lower half of the normal probability paper and relabel the upper part as P' where

$$P' = 2P - 1.$$

A rough ruling of this graph is given in Figure 1.

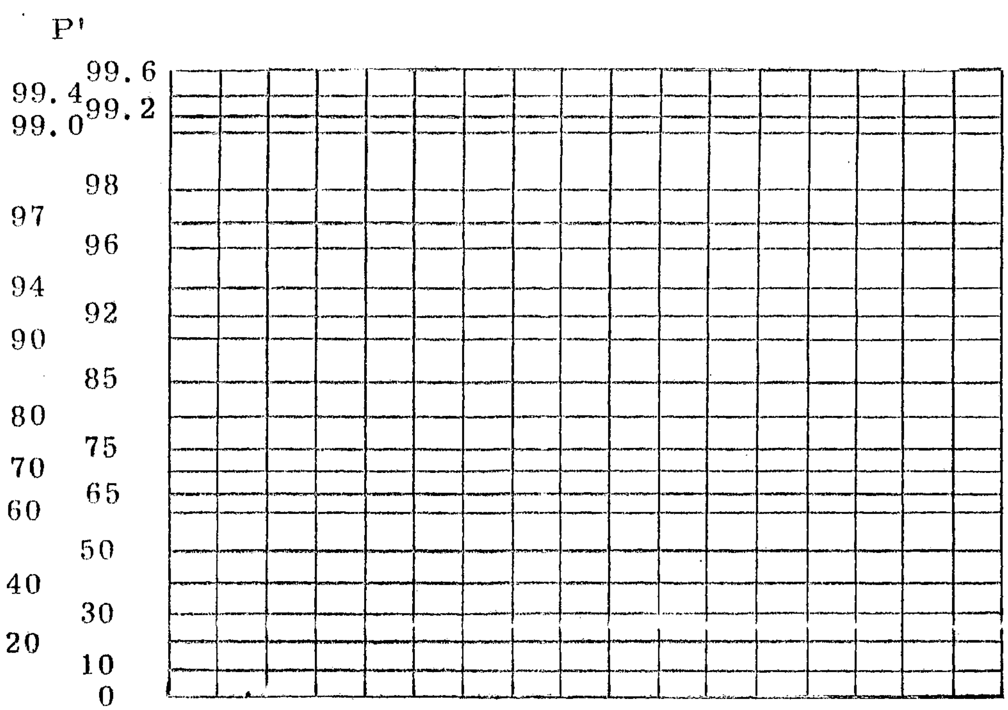

Fig. 1 - General purpose half-normal grid.

To plot the empirical distribution we take the absolute values of the contrasts and order them from the smallest to the largest. We have, of course, a discrete set of numbers and we take the empirical probability of the i-th largest to be

$$P'_i = (i - \tfrac{1}{2})/n$$

where n is the total number of contrasts. The abscissa corresponding to the empirical probability P'_i is the value of the i-th largest-in-absolute-value contrast.

Cuthbert Daniel[5] has constructed special purpose grids for four of the most common 2^p factorial experiments. They are the 2^4, 2^5, 2^6, and 2^7 factorial experiments where $n = 15, 31, 63$, and 127 respectively. These are simpler to use since no probabilities need to be computed and the rank numbers are printed directly on the grids. These four special purpose grids are shown in Figure 2. The largest contrast is plotted on the top line and so on down to the smallest. However, there is little advantage in plotting each contrast of the smaller half of the set because of the enforced close correlation in magnitude between adjacent ordered values.

After the contrasts have been plotted, or even before, we can construct the straight line through the origin and with slope $1/\sigma_c$ with which we wish to compare our contrasts. Of course, we do not know σ_c, but we have several ways we can estimate it. One estimate which is very easy to obtain is to use the value of the contrast for which P' is most nearly 0.683. This is a logical choice for an estimate of σ_c since if this is a half-normal population we are working with, we can expect 68.3% of the population to be between zero and σ_c. For values of n of 15, 31, 63, and 127, the corresponding σ_c estimators are the 11th, 22nd, 44th, and 88th ordered contrast.

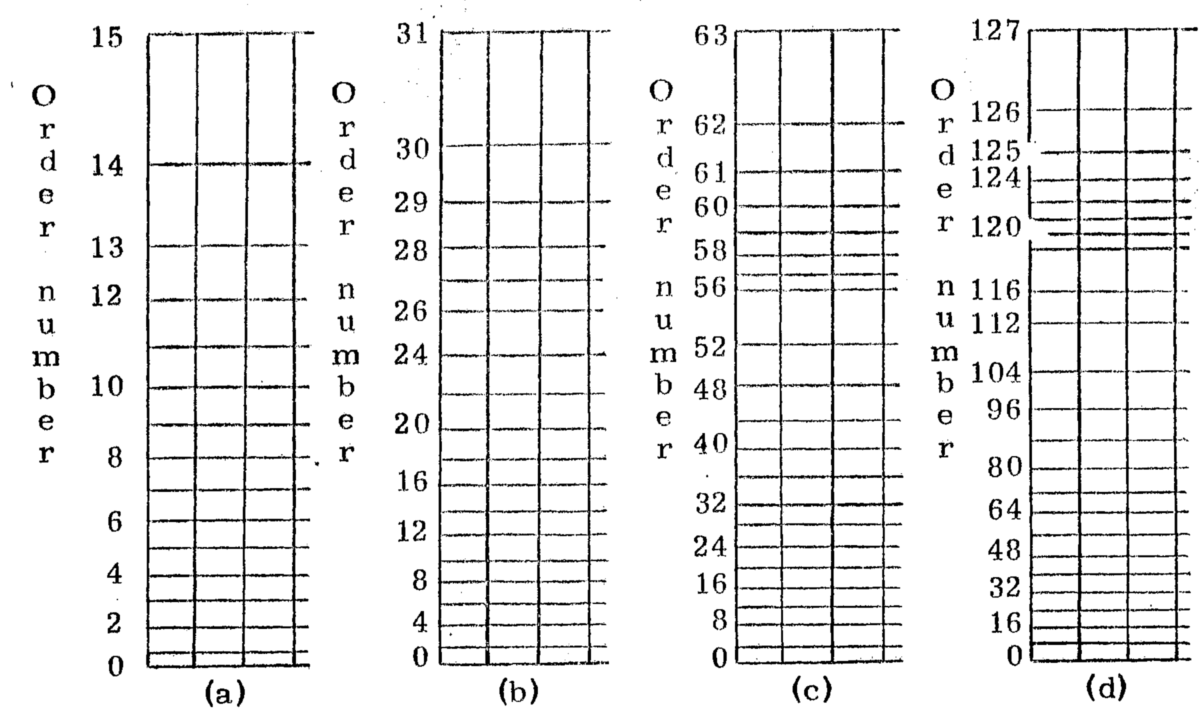

Fig. 2 - Half-normal scales for (a) 15 d.f., (b) 31 d.f., (c) 63 d.f., (d) 127 d.f.

To get some idea of the sampling variation inherent in this form of plotting, Cuthbert Daniel[5] has plotted ten sets of 31 random standard normal deviates taken from Dixon and Massey's tables[7]. The first 31 in each of the columns 11 to 21 are used. The printed abscissa values are correct only for the first plot in each set. Each sloping line should start at zero. Each has the slope required by the population standard deviation (1.00) and not the slope developed in the previous paragraph. These ten plots are displayed in Figure 3.

After the contrasts have been plotted and the half-normal line has been constructed, we can look at the finished plot to decide if the contrasts plotted are, in fact, all from the same half-normal population, the population of error contrasts. Any large deviation of a contrast to the right of the half-normal line would indicate that the contrast was probably not from the population of error contrasts. We can then drop all such contrasts and plot the remaining as we did before. Of course, all of the P' values would change as well as our estimate of σ_c. If we have removed all of the contrasts that are not error contrasts, the new plot should have no large deviations. We can now use these contrasts to estimate error. The question of when a deviation become large enough to be removed will be discussed later.

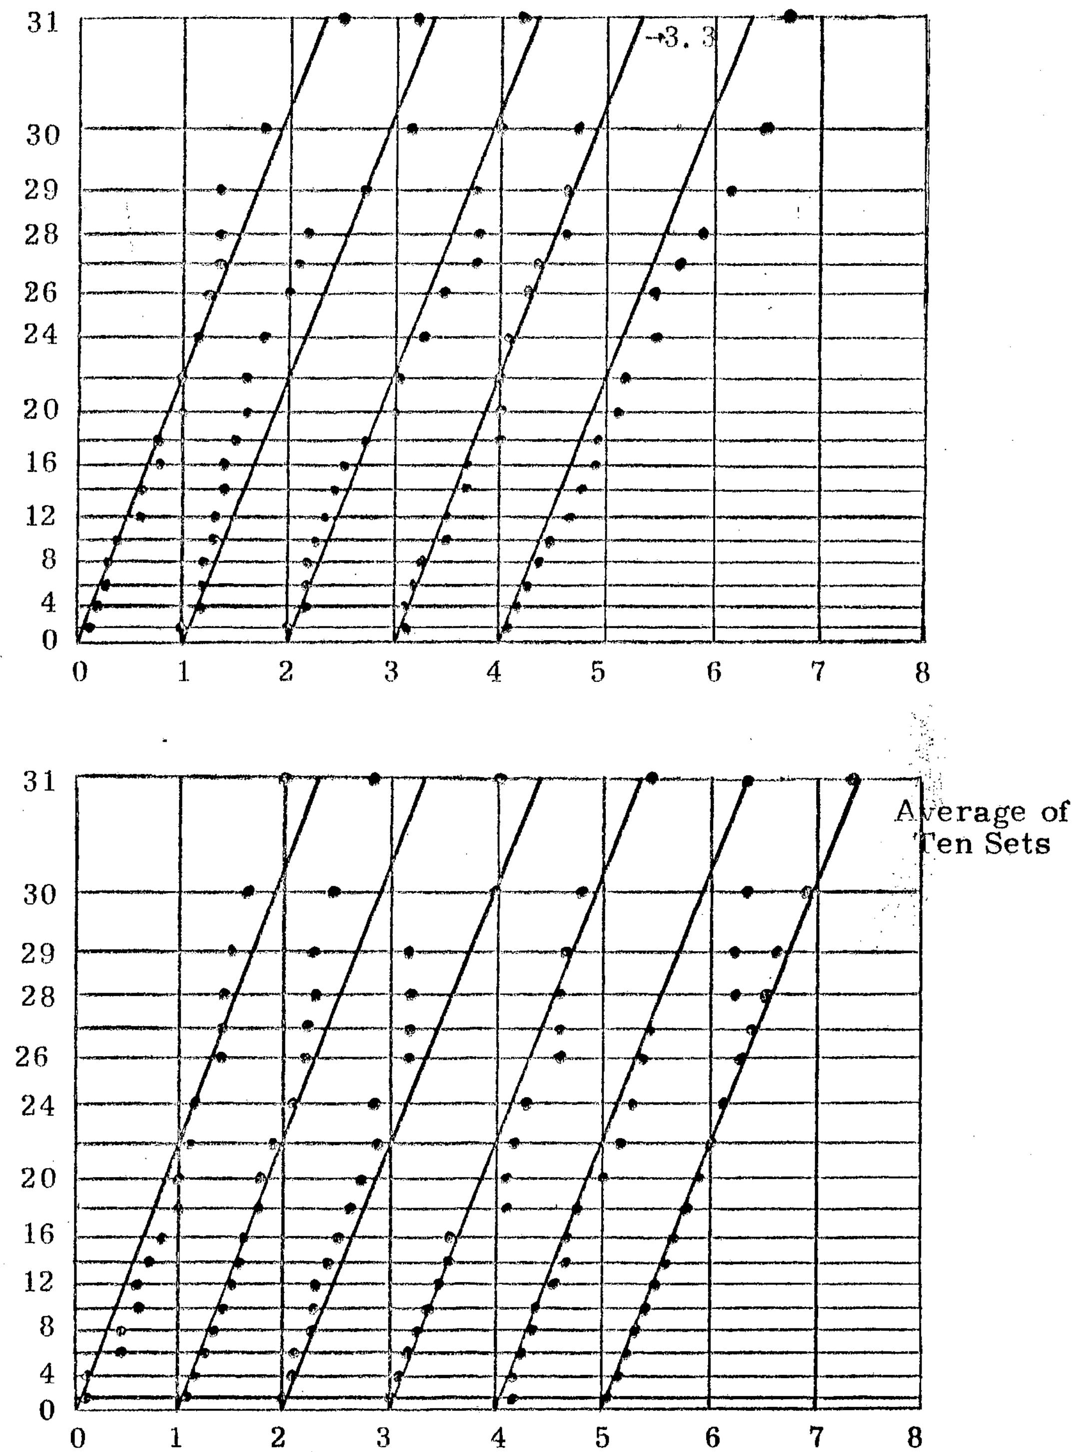

Fig. 3 - Half-normal plots of ten sets of 31 random standard normal deviates.

CHAPTER III

AN ALTERNATE BASIC ASSUMPTION FOR THE ANALYSIS OF TWO-LEVEL FACTORIAL EXPERIMENTS

In a 2^p factorial experiment without replication, there is no estimate of error if all effects and interactions are real. The standard inference methods require the assumption that certain of the main effects or interactions are not real. The sums of squares due to these main effects and interactions are then used to estimate the error sum of squares. It is a common practice to assign all interactions above a certain number of factors to error, since the higher order interactions are the ones most likely to be estimates of error.

Experience in certain areas of experimentation suggests that usually the number of true main effects or interactions will be quite small, relative to n, and in particular, that it will be much smaller than the number of main effects and interactions that were assumed to be real. Experience also suggests that higher order interactions are sometimes actually real.

The preceding points suggest that maybe an alternative to the standard assumption of null higher order interactions should be considered. They also suggest that as a basis for inference, a new underlying assumption might be that any of the n main effects or interactions may be real but that, at most, only a certain small number of them are real. The statistical problem, then, is to infer which, if any, of the n main effects or interactions are real.

This problem adapts itself very well to the half-normal plotting procedure. In our half-normal plotting procedure, all of the contrasts that appear *not* to be from the population of error contrasts are removed and a new half-normal plot is constructed. This is repeated until it appears that we have eliminated all of the contrasts that do not belong to the population of error contrasts. In other words, we remove all of the contrasts that are real and we consider the main effects or interactions corresponding to these contrasts as real also.

Davies [6] discusses a 2^5 experiment on penicillin production which is analyzed below under these new assumptions. The 31 contrasts are arranged in decreasing order of absolute magnitude in Table I. For convenience, the values have been multiplied by 100 and rounded.

TABLE I

No.	Effect	Value	No.	Effect	Value
31	E	224	15	DE	30
30	A	190	14	BE	29
29	C	153	13	BDE	28
28	CE	93	12	ABE	22
27	ABCDE	77	11	ADE	21
26	AB	64	10	BCD	18
25	ABCD	58	9	BCDE	16
24	ACE	58	8	ABDE	14
23	AD	54	7	CDE	12
22	AC	53	6	D	9
21	BC	53	5	BD	7
20	ACDE	47	4	B	6
19	BCE	39	3	CD	4
18	ABD	34	2	AE	2
17	ACD	33	1	ABC	0
16	ABCE	31			

Figure 4 is the plot of this data on a special half-normal grid where n equals 31. Since the value of the 22nd ordered contrast is 53, our estimate of σ_c is 53 and we construct the half-normal line through the origin and with slope 1/53. Actually, all we do is to draw a straight line through the origin and the point on the graph where the value of the 22nd ordered contrast is plotted. It is clear from looking at the plot that E, A, and C are probably not error contrasts. We are told that CE was judged likely to be appreciable from previous information. If we remove these four contrasts and re-plot the remaining 27 contrasts, we get the dashed line in Figure 4 where the x's are the values of the largest of the 27 contrasts. Our estimate of σ_c is now the value of the 19th ordered contrast or 39. It appears from looking at this new plot that we have eliminated all contrasts that are not error contrasts. We can then use these 27 main effects and interactions to estimate error and we conclude that E, A, C, and CE are probably real. This agrees with the conclusion reached by those who ran this experiment except they believed that ABCDE should be significiant also since it was confounded with blocks.

Another example from Davies[6] is a 2^4 experiment on the preparation of an isatin derivative from an isonitrosacetylamine derivative. Again, for convenience, the data has been multiplied by 100 and rounded. The 15 contrasts are arranged in decreasing order of absolute magnitude in Table II.

Figure 5 is the plot of this data on a special half-normal grid where n equals 15. Since the value of the 11th ordered contrast is 15, our estimate of σ_c is 15 and we construct the half-normal line through the origin and with slope 1/15. Since there are no large

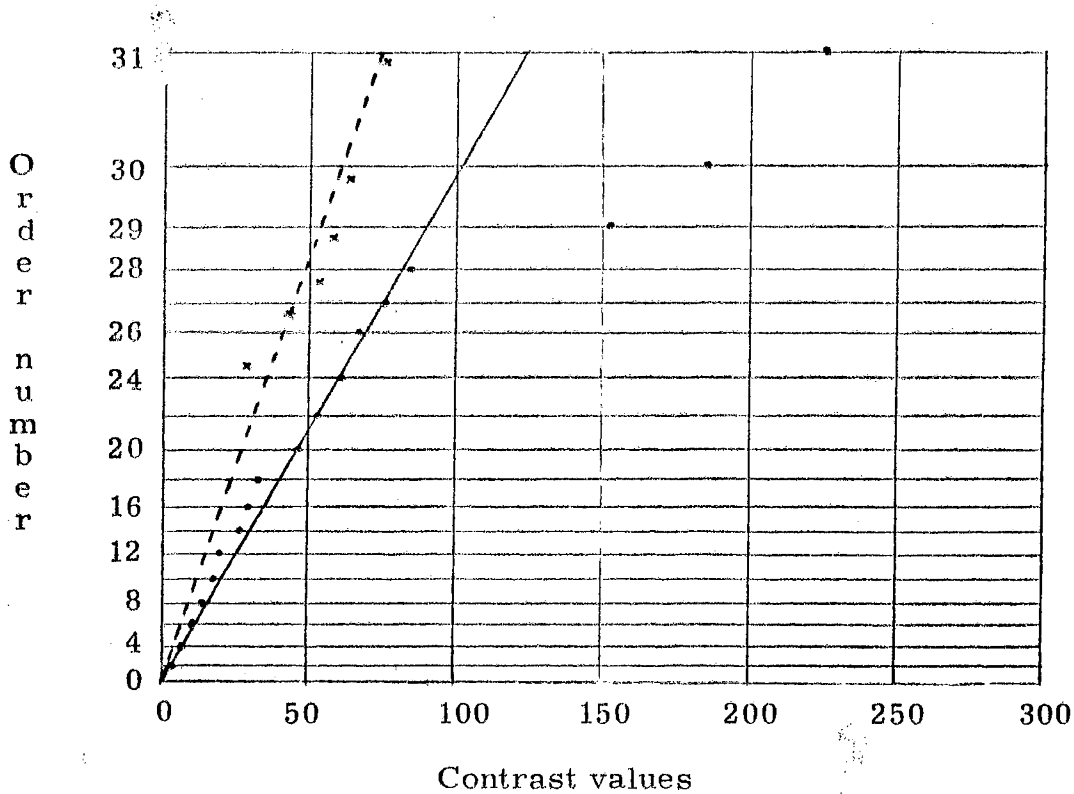

Fig. 4 - Half-normal plot of a 2^5 experiment

TABLE II

No.	Effect	Value	No.	Effect	Value
15	D	27	7	BC	6
14	BD	25	6	AC	3
13	A	19	5	CD	3
12	AD	16	4	B	2
11	ABC	15	3	ABCD	2
10	BCD	12	2	ACD	1
9	ABD	10	1	AB	0
8	C	8			

deviations to the right of the half-normal line, it appears that all of the contrasts are from the population of error contrasts and that there are no true main effects or interactions in the experiment. This does not agree with the conclusion reached by those who ran the experiment. They used the standard inference method and reached the conclusion that D was significant at the five per cent level and that BD, while not quite significant at the five percent level, should be considered significant because of prior knowledge about this interaction.

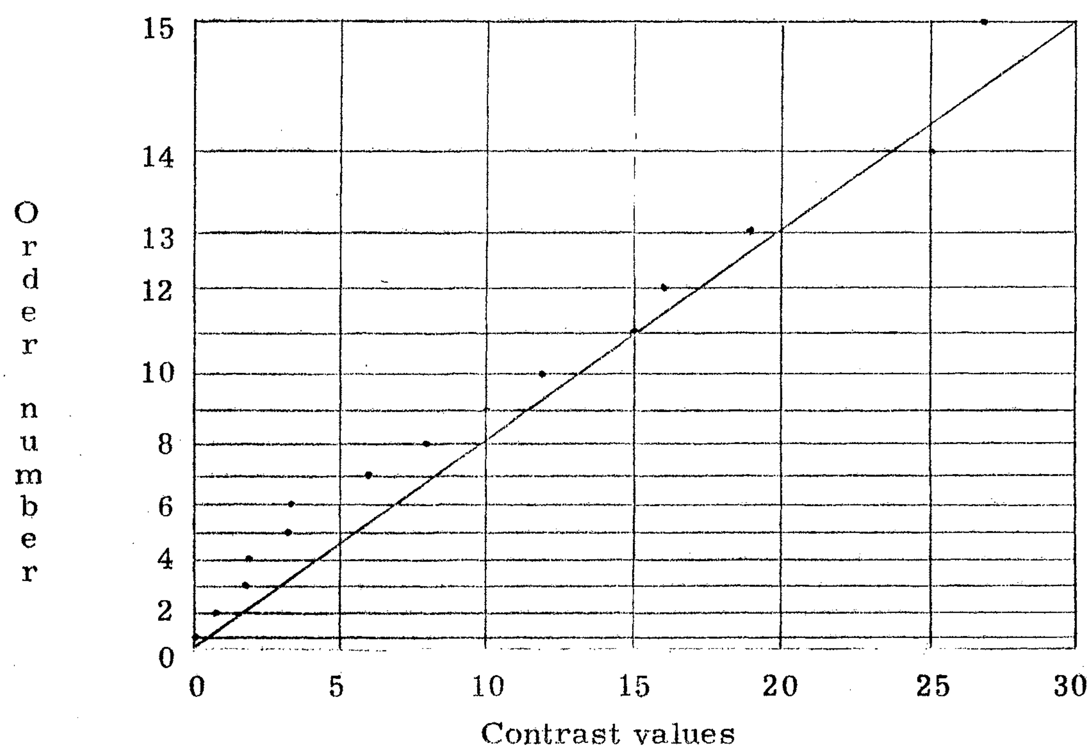

Fig. 5 - Half-normal plot of a 2^4 experiment

A possible explanation for this difference in the conclusions reached might be that if this is a null experiment, the probability that at least one effect will show significance is not .05 but $1 - .95^{10} = 1 - .60 = .40$. This results from the fact that in a null experiment, each of the effects tested, in this case ten, have the same probability of being considered significant. Since we are dealing in this case with a Type I error rate of .05, the probability that none of the effects will appear significant is $.95^{10}$ and the probability that at least one effect will appear significant is $1 - .95^{10}$. If we compare Figure 5 with Figure 11, we see that the D effect is near to the $\alpha = .40$ line.

CHAPTER IV

A TEST STATISTIC FOR HALF-NORMAL PLOTS

Allen Birnbaum[1] has investigated the usefulness of half-normal plots in distinguishing between the cases of no effects and those of one effect. The distributions of interest are those of

$$t_n^{(\Delta)} = \frac{u_n}{u_a}$$

where:

u_n is the largest in absolute magnitude of n observed values of a random normal variable with population mean zero, variance σ^2, but with one value chosen at random changed by the addition to its signed value of $\Delta\sigma$;

u_a is the absolute value of the order-statistic from the same set as u_n that is numbered nearest to $(0.683n + 0.5)$ and which therefore most nearly estimates σ directly in a null experiment. For values of n of 15, 31, 63, and 127, the corresponding σ-estimators are the 11th, 22nd, 44th, and 88th order statistic;

$t_n^{(\Delta)}$ is then the ratio of the largest in absolute value of the n observed values (one perturbed by the addition of $\Delta\sigma$) to the corresponding σ-estimator, u_a.

Birnbaum has studied the distribution of $t_n^{(\Delta)}$ for n = 31, 63, and 127 and for Δ's of 0 to 6. At his request, a machine sampling of 2500 sets of 31 random standard normal deviates was carried out by G. L. Lieberman and associates at Stanford University. This information permitted an estimate of the distribution of $t_{31}^{(0)}$ with good accuracy over the range of interest. It was found that the distribution of $\log_{10} t_{31}^{(0)}$ is quite closely approximated by the normal distribution with mean 0. 35 and standard deviation 0. 11. The points on Figure 6 give several values from the empirical distribution of logarithms plotted on a normal probability scale. Daniel[5] reports that H. M. Truax under slightly different assumptions found values closely approximating those given above.

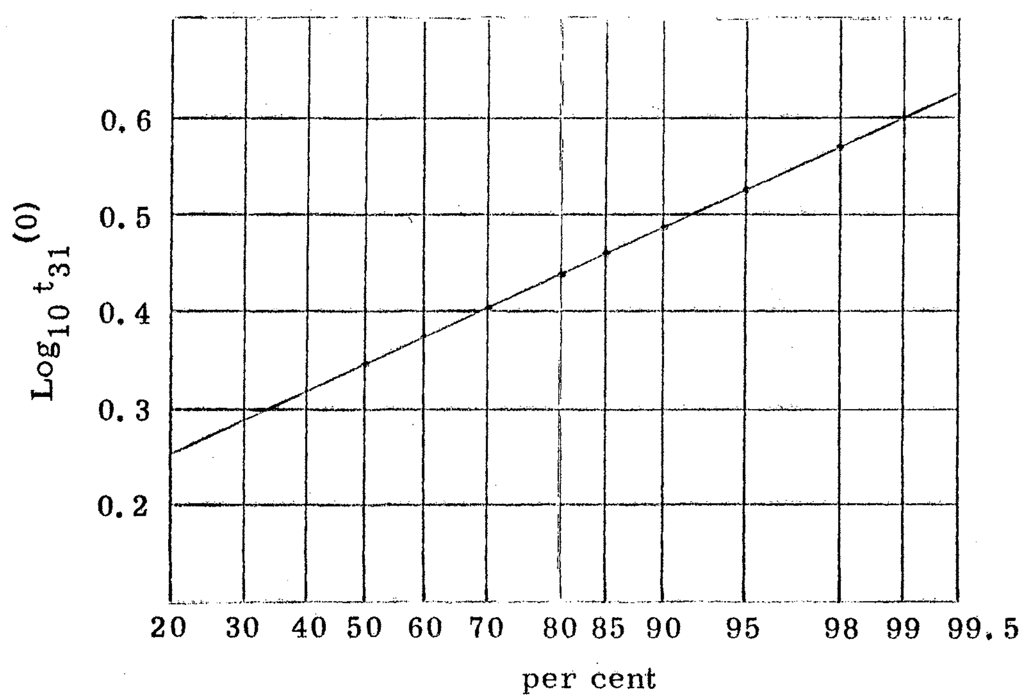

Fig. 6 - Empirical cumulative distribution of $\log_{10} t_{31}^{(0)}$

Birnbaum's paper gives (asymptotic) approximations for the distributions of $t_n^{(0)}$ and uses them to approximate the distributions for n equal to 63 and 127. Daniel [5] has estimated the distribution of $t_{15}^{(0)}$ using 198 samples of 15 random standard normal deviates taken from the tables of Dixon and Massey[7]. The approximate distribution of $\log_{10} t_n^{(0)}$ for n equal to 15, 31, 63 and 127 is given in Table III and the estimated cumulative distributions are plotted in Figure 7.

TABLE III

Approximate Distribution of $\log_{10} t_n^{(0)}$

All are nearly normal for $P > 0.1$ with estimated parameters $\hat{\mu}$, $\hat{\sigma}$.

n	$\hat{\mu}$	$\hat{\sigma}$
15	0.265	0.135
31	0.345	0.115
63	0.405	0.078
127	0.425	0.078

Let $y_1, y_2, \cdots, y_n$ be the n unordered constrasts and let k be a suitable positive constant such that if $t_n \leq k$, it is inferred that all of the contrasts are from the population of null contrasts. If $t_n > k$, it is inferred that one of the contrasts, the one corresponding to u_n, is not from the population of null contrasts.

In order to study the probabilities of the various possible types of errors associated with the use of the above described procedures based on t_n, the distribution of t_n must be considered under each of the following hypothesis:

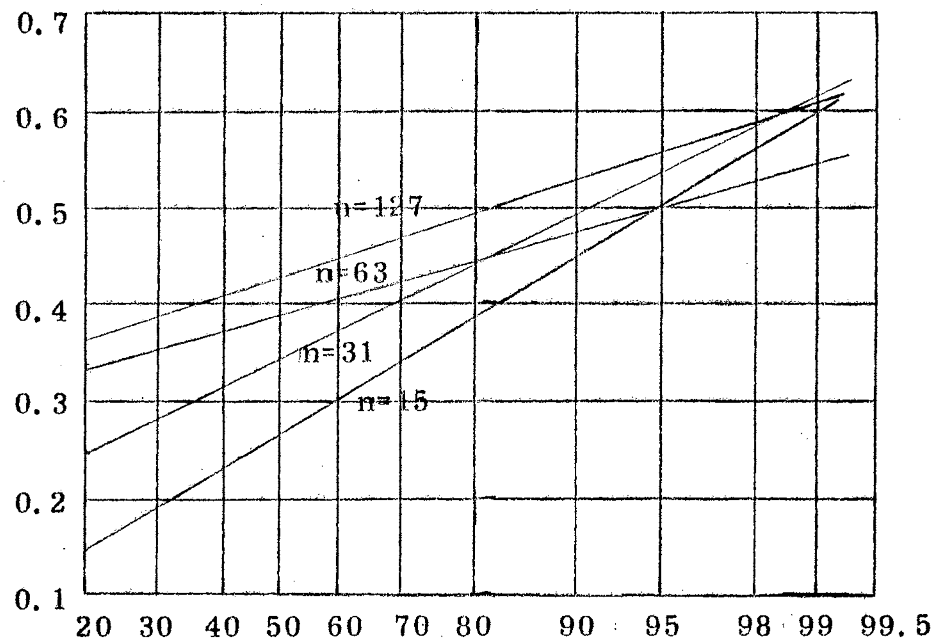

Fig. 7 - Estimated cumulative distribution of $\log_{10} t_n^{(0)}$ for $n = 15, 31, 63, 127.$

H_o: $E(y_i) = 0, \; i = 1, 2, \cdots, n. \quad \Delta = 0$

H_Δ; $E(y_j) = \Delta\sigma$ for just one unknown value j, while

$E(y_i) = 0$ for each $i \neq j$. Δ is a non-zero constant.

In order to test H_o against H_Δ at a given significance level α, it is required that

$$1 - \alpha = \Pr(t_n \leq k \mid H_o)$$

where $k = k(n, \alpha)$.

Where n is large, $u_a \doteq \sigma$ with probability very near to one under both H_o and H_Δ, for any Δ. Also, t_n is distributed almost identically with u_n/σ under both H_o and H_Δ. Therefore,

$$1 - \alpha = \Pr(t_n \leq k \mid H_o) = \Pr(u_n/u_a \leq k|H_o)$$

$$\doteq \Pr(u_n/\sigma \leq k|H_o) = \Pr(|y_i|/\sigma \leq k|H_o \text{ for } i = 1, 2, \cdots, n)$$

$$= \Pr(|y_1|/\sigma < k|H_o) \Pr(|y_2|/\sigma < k|H_o) \cdots \Pr(|y_n|/\sigma < k|H_o)$$

$$= \left[\int_{-k}^{k} \frac{1}{\sqrt{2\pi}} e^{\frac{-x^2}{2}} dx\right]^n = \left[2\int_0^k \frac{1}{\sqrt{2\pi}} e^{\frac{-x^2}{2}} dx\right]^n$$

$$= \left[2\int_{-\infty}^{k} \frac{1}{\sqrt{2\pi}} e^{\frac{-x^2}{2}} dx - 2\int_{-\infty}^{0} \frac{1}{\sqrt{2\pi}} e^{\frac{-x^2}{2}} dx\right]^n = [2\phi(k) - 1]^n$$

where

$$\phi(k) = \int_{-\infty}^{k} \frac{1}{\sqrt{2\pi}} e^{\frac{-x^2}{2}} dx$$

Let

$$k^* = k^*(n, \alpha)$$

such that

$$1 - \alpha = [2\phi(k^*) - 1]^n$$

$$(1 - \alpha)^{\frac{1}{n}} = 2\phi(k^*) - 1$$

$$2\phi(k^*) = 1 + (1 - \alpha)^{\frac{1}{n}}$$

$$\phi(k^*) = \frac{1}{2} + \frac{(1 - \alpha)^{\frac{1}{n}}}{2}$$

$$k^* = \phi^{-1}\left(\frac{1}{2} + \frac{(1 - \alpha)^{\frac{1}{n}}}{2}\right)$$

The effect of approximating $k(n, \alpha)$ by $k^*(n, \alpha)$ is represented by the discrepancy

$$D(n, \alpha) = \Pr(t_n < k^*|H_o) - \alpha$$

$D(n, \alpha)$ is the increase in Type I error rate caused by using k* instead of k as a critical value. Table IV gives some values of $D(n, \alpha)$ for n equals 31 as calculated by Allan Birnbaum[1]. It is seen that $D(31, \alpha) < .035$ for $\alpha \leq .5$. Since $D(n, \alpha)$ decreases to zero as n increases for any fixed α, the values given provide a useful indication that for $n > 31$, the approximation of k by k* is satisfactory for most practical purposes.

TABLE IV

α	$k(31, \alpha)$	$D(31, \alpha)$
.01	3.93	.02
.02	3.75	.02
.05	3.36	.03
.10	3.06	.03
.20	2.75	.02
.30	2.54	.02
.40	2.38	.01
.50	2.24	.03

The values of k* can now be computed from any full tables of ϕ or ϕ^{-1}. Some values of k* where n equals 63 and 127 and α equals .40, .20, .10, .05, and .01 as found by Birnbaum[1] are presented in Table V

TABLE V

α	n = 63	n = 127
.01	3.78	3.95
.05	3.35	3.54
.10	3.14	3.34
.20	2.92	3.13
.40	2.65	2.88

Under H_Δ, for a large n the probability of Type II error is

$$\Pr(t_n \le k|H_\Delta) = \Pr(u_n/u_a \le k|H_\Delta) \doteq \Pr(u_n/\sigma \le k|H_\Delta)$$

$$= \Pr(|y_i|/\sigma \le k|H_\Delta, \text{ for } i = 1, 2, \cdots, n)$$

$$= \Pr(|y_1|/\sigma \le k|H_\Delta)\Pr(|y_2|/\sigma \le k|H_\Delta)\cdots\Pr(|y_n|/\sigma \le k|H_\Delta)$$

$$= \left[\int_{-k}^{k} \frac{1}{\sqrt{2\pi}} e^{\frac{-x^2}{2}} dx\right]^{n-1} \left[\int_{-k-\Delta}^{k-\Delta} \frac{1}{\sqrt{2\pi}} e^{\frac{-x^2}{2}} dx\right]$$

$$= (1-\alpha)^{\frac{n-1}{n}} \left[\phi(k-\Delta) - \phi(-k-\Delta)\right]$$

$$\doteq (1-\alpha)\left[\phi(k-\Delta) - \phi(-k-\Delta)\right] .$$

Thus with tables of ϕ a simple calculation gives approximately the power function

$$1 - \Pr(\text{type II error}) = 1 - (1-\alpha)\left[\phi(k-\Delta) - \phi(-k-\Delta)\right]$$

A related operating characteristic of interest is the sensitivity function

$$\gamma = \gamma(n,\alpha,\Delta) = \Pr(|y_i|/u_a > k|\Delta)$$

$$= 1 - \Pr(|y_i|/u_a \le k|\Delta) = 1 - \int_{-k-\Delta}^{k-\Delta} \frac{1}{\sqrt{2\pi}} e^{\frac{-x^2}{2}} dx$$

$$= \int_{-\infty}^{-k-\Delta} \frac{1}{\sqrt{2\pi}} e^{\frac{-x^2}{2}} dx + \int_{k-\Delta}^{\infty} \frac{1}{\sqrt{2\pi}} e^{\frac{-x^2}{2}} dx$$

$$= \int_{-\infty}^{-k-\Delta} \frac{1}{\sqrt{2\pi}} e^{\frac{-x^2}{2}} dx + \int_{-\infty}^{-k+\Delta} \frac{1}{\sqrt{2\pi}} e^{\frac{-x^2}{2}} dx$$

$$= \phi(-k-\Delta) + \phi(-k+\Delta) .$$

γ may be interpreted as the probability that the correct observation y_i, where $E(y_i) = \sigma\Delta \neq 0$, will appear as significant, that is that $|y_i|/u_a > k$ whether or not $|y_i| = u_m$. For a large n and with $\Delta = k$

$$\gamma(n, \alpha, \Delta) \doteq \tfrac{1}{2} + \phi(-2k) \text{ and the power function}$$

$$1 - \beta \doteq 1 - (1-\alpha)\left[\tfrac{1}{2} - \phi(-2k)\right] .$$

If $\Delta = k \geq 1.3$, then

$$\gamma \doteq \tfrac{1}{2} \text{ and } 1 - \beta \doteq \tfrac{1}{2} + \alpha/2.$$

Thus to obtain sensitivity $\gamma = \frac{1}{2}$ against any specified value of $\Delta \geq 1.3$, let the critical value $k = \Delta$, which gives a Type I error rate of

$$\alpha \doteq 1 - \left[2\phi(\Delta) - 1\right]^n .$$

The upper and lower bounds on the sensitivity function as found by Birnbaum[1] for n equals 31 and α equals .05, .20 and .40 are shown in Figure 8 by the three pair of lines that converge to the right. Empirical sampling by Daniel[5] for Δ equals 4 using 99 sets of 31 random standard normal deviates gives the values shown by the three dots in that figure.

A similar graph for n equals 15, estimated by Daniel[5] from a sample of 198 sets of 15 random standard normal deviates is given in Figure 9. Since Δ's of 2 and 4 only were used, the straight lines drawn are entirely conjectural, based on the analogy with the linear bounds found by Birnbaum for n equals 31.

It must be pointed out that the preceding discussion is based on the strong assumption that at most one of the main effects or interactions is real. If many of them are real, the power and sensitivity of tests based on t_n may be much reduced. However, if only a small

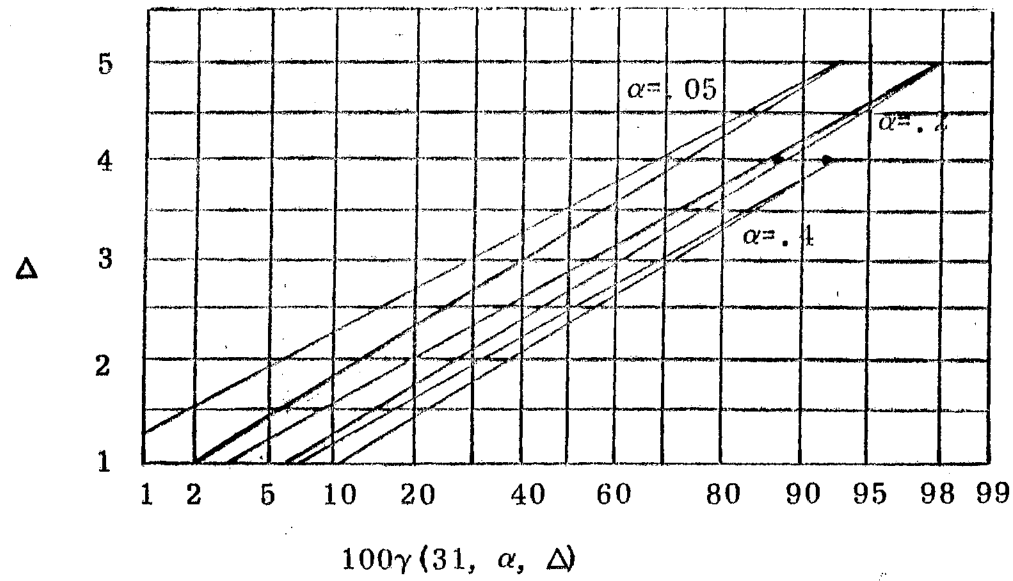

Fig. 8 - Bounds on sensitivity function of $t_{31}^{(\Delta)}$ for one effect of size Δ.

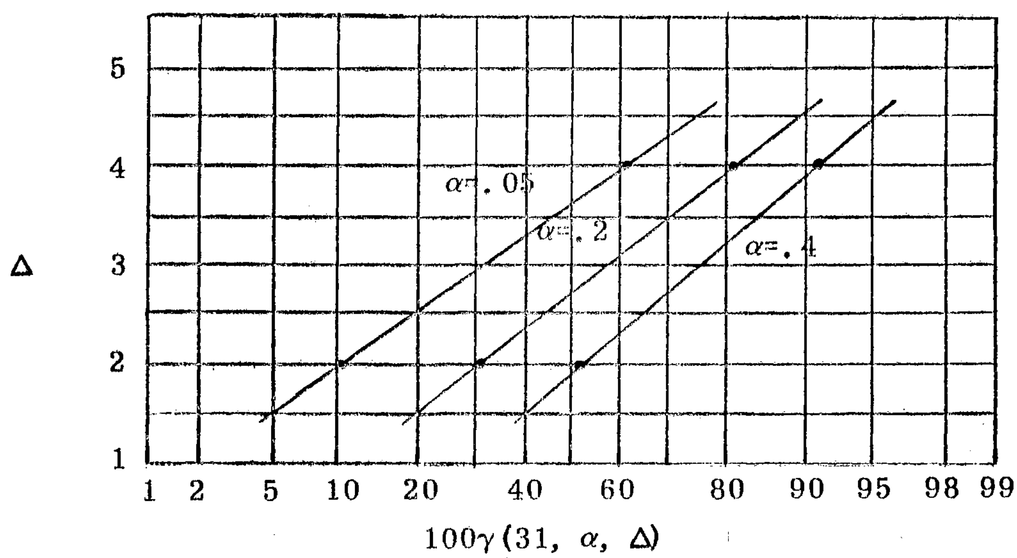

Fig. 9 - Approximate sensitivity function of $t_{15}^{(\Delta)}$ for one effect (based on 198 sets of random normal deviates).

number of the main effects or interactions is real, it seems likely that the power and sensitivity properties indicated above will tend to hold approximately, with appropriate reinterpretation for the various possible alternative hypothesis.

Since more than one non-zero effect or interaction is usually expected, the distributions of $t_{n-1}^{(0)}$, $t_{n-2}^{(0)}$, etc. are also of interest. Daniel[5] has estimated the cumulative distribution of $t_{n-1}^{(0)}$ and $t_{n-2}^{(0)}$ empirically using the same 198 sets of 15 random standard normal deviates. Figure 10 shows the empirical cumulative distribution of $\log_{10} t_{15-j}^{(0)}$ for j = 0, 1, and 2.

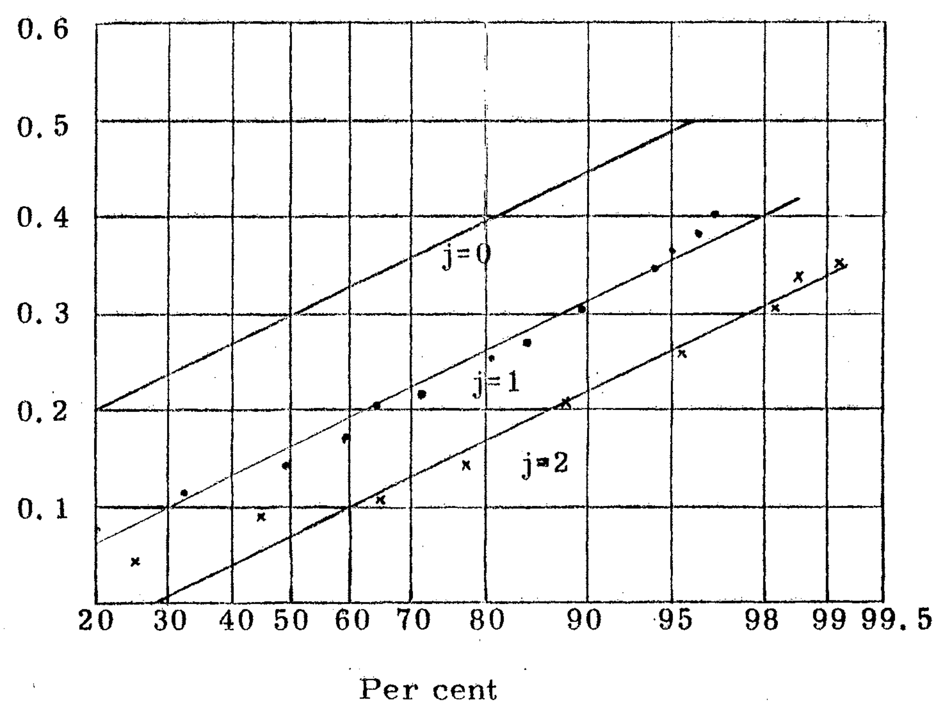

Fig. 10 - Empirical cumulative distribution for $\log_{10} t_{15-j}^{(0)}$ for j = 0, 1,

CHAPTER V

STANDARDIZED HALF-NORMAL PLOTS

A half-normal plot on which the half-normal line and the probability of Type I error are already plotted would be convenient in the study of the results of 2^p factorial experiments. Dividing the computed ranked contrasts by the σ-estimator contrast will give a scale-free set of order-statistics that should in the absence of any real effects fall around a half-normal line plotted through the origin and a point on the σ-estimator line with abscissa of one. The desired probability of Type I error, α, can now be chosen. Using Figure 7, the point on the estimated cumulative distribution of $\log_{10} t_n^{(0)}$ line where it crosses the $1 - \alpha$ line is found. This gives a number whose antilogarithm can easily be found. This number is now plotted on the top line of the half-normal plot. The probability that a null contrast will fall to the right of this point is α.

As an example of how this works, for $n = 15$ and $\alpha = .05$, the $n = 15$ line crosses the $1 - .05$ or $.95$ line at the point $.49$. This is a number whose antilogarithm is 3.09. This number, 3.09, is now plotted on the line marked 15. This now gives that the probability, in a null experiment, that $u_{15}/u_{11} > 3.09$ is 5% or that the probability that $u'_{15} = u_{15}/u_{11}$ when plotted on line 15 will fall to the right of 3.09 is 5%.

Another method which can be used to find this critical number, and which is probably more accurate, is to use the information in Table III to solve the following equation for the desired values of n and α.

$$\int_{-\infty}^{Z} \frac{1}{\sqrt{2\pi}\hat{\sigma}} \exp - \frac{1}{2}\left(\frac{x - \hat{\mu}}{\hat{\sigma}}\right)^2 \quad dx = 1 - \alpha$$

which under the transformation $y = (x - \mu)/\sigma$ becomes

$$\int_{-\infty}^{\frac{Z-\hat{\mu}}{\hat{\sigma}}} \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{y^2}{2}\right) dy = 1 - \alpha .$$

By using the cumulative normal distribution tables, the value of $(Z - \hat{\mu})/\hat{\sigma}$ can be found for any value of $1 - \alpha$. After substituting the values of $\hat{\mu}$ and $\hat{\sigma}$, Z and its antilogarithm can be found. It is this value, the antilogarithm of Z, that is plotted on the top line of the half-normal graph. The probability that a null contrast will fall to the right of this point is α. Table VI, which was developed by this second method, gives the critical values for the usual four values of n and for Type I error rates of .01, .05, .10, .20, and .40.

TABLE VI

Critical Values for Standardized Half-normal Plots

n	0.01	0.05	0.10	0.20	0.40
15	3.79	3.07	2.74	2.39	1.92
31	4.10	3.42	3.11	2.77	2.37
63	3.86	3.41	3.20	2.96	2.66
127	4.04	3.58	3.35	3.10	2.79

Figure 11 is the standardized half-normal plot for n equals 15. The points on line 15 are taken from Table VI while the dotted lines for α

equal to .05, .20 and .40 are as deduced by Daniel[5] from Figure 10. Figures 12, 13, and 14 are the standardized half-normal plots for n equals 31, 63, and 127 respectively. The points on the top lines are taken from Table VI while the dashed lines are contributed by Daniel[5].

As an example of the use of a standardized half-normal grid, we take another 2^4 experiment from Davies[6]. Table VII gives the original values as well as the scale-free values, u_i/u_{11}.

TABLE VII

No.	Value	u_i/u_{11}	No.	Value	u_i/u_{11}
15	41.91	3.81	7	5.28	.48
14	16.72	1.52	6	4.73	.43
13	13.64	1.24	5	3.63	.33
12	12.98	1.18	4	2.64	.24
11	11.00	1.00	3	1.43	.13
10	8.36	.76	2	1.21	.11
9	7.92	.72	1	.77	.07
8	7.15	.65			

When the data from the u_i/u_{11} column from Table VII are plotted on the standardized half-normal plot for $n = 15$, the results are as in Figure 15. This figure shows that the 15th ordered contrast is significant at around the one per cent level and that the 14th and 13th ordered contrasts are significant at around the forty per cent level.

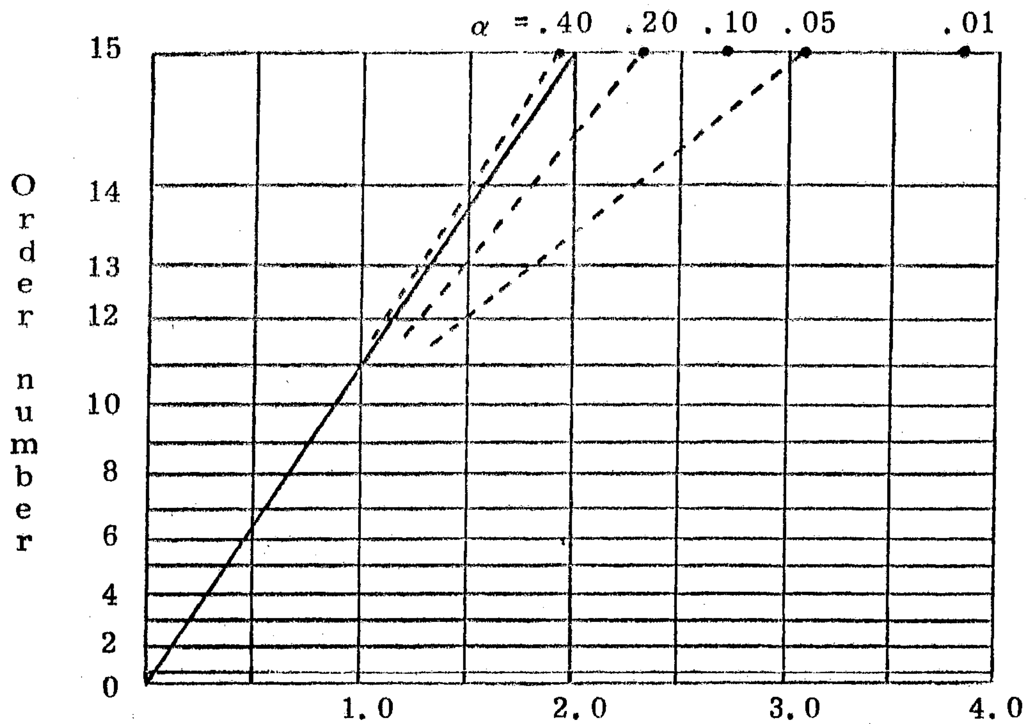

Fig. 11 - Standardized half-normal grid for 15 contrasts

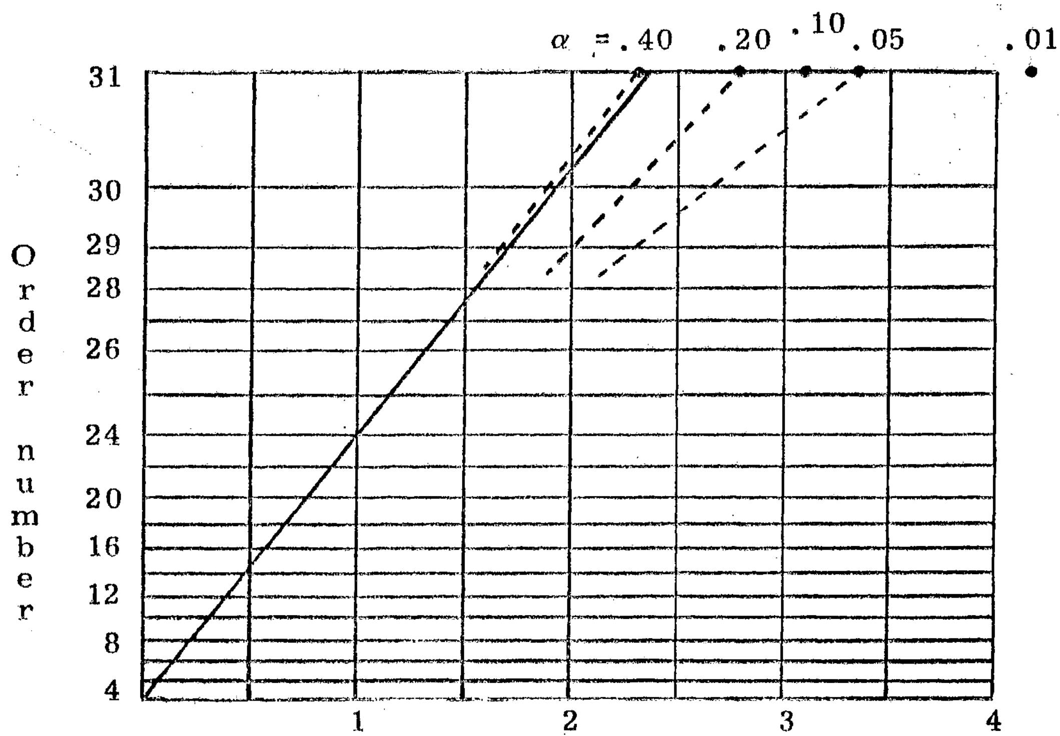

Fig. 12 - Standardized half-normal grid for 31 contrasts

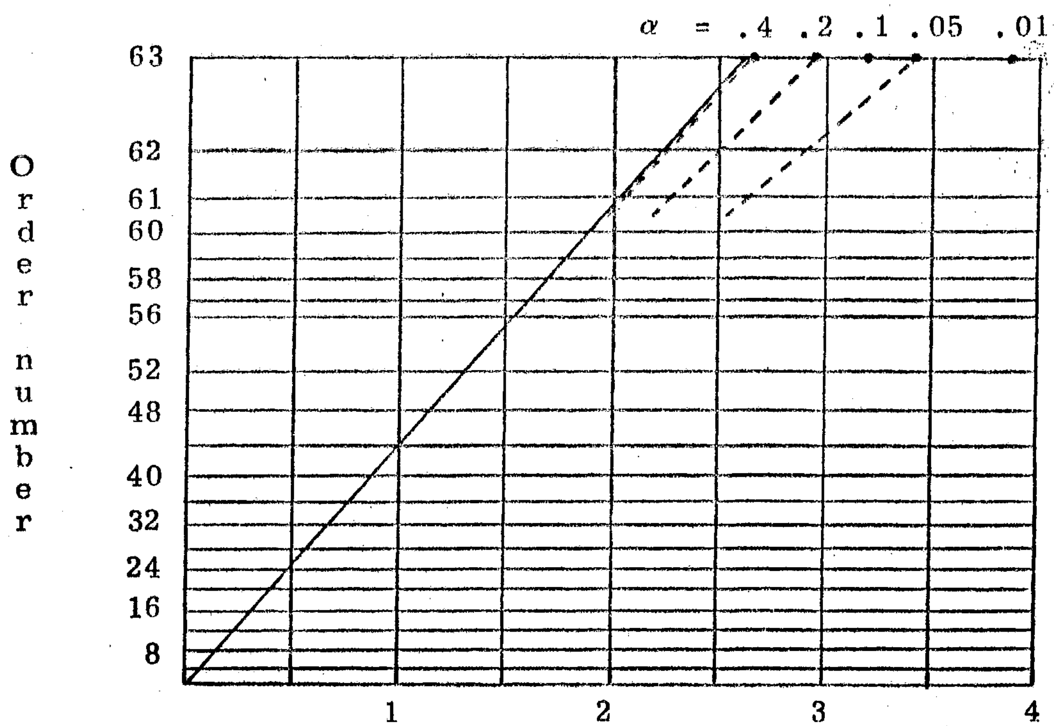

Fig. 13 - Standardized half-normal grid for 63 contrasts

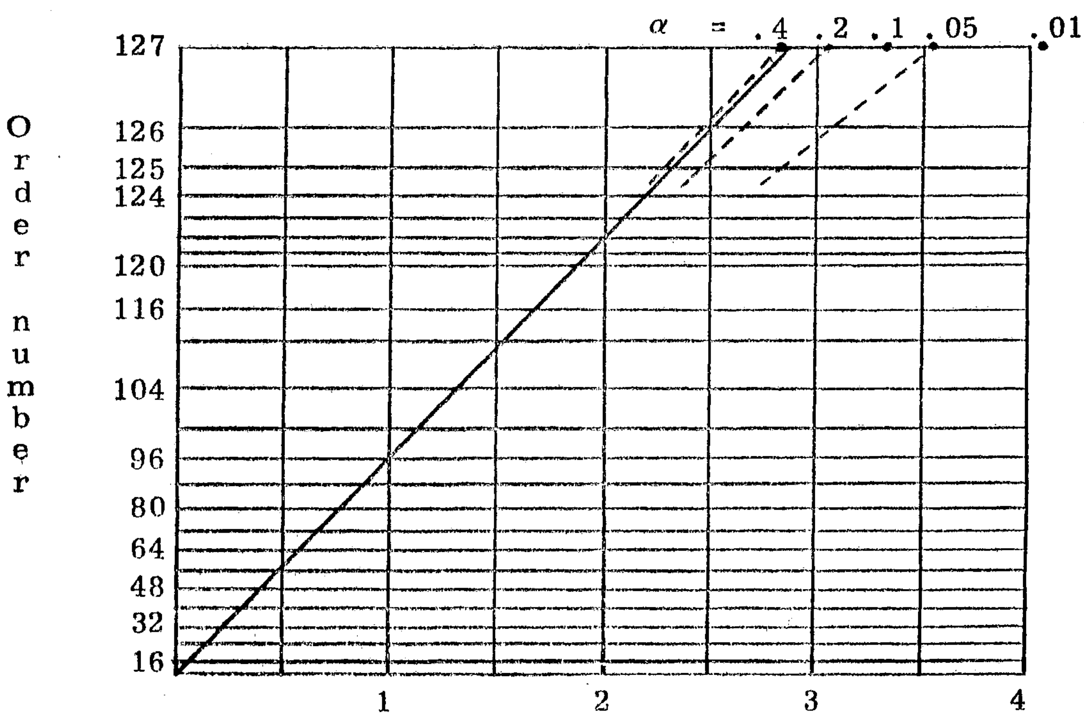

Fig. 14 - Standardized half-normal grid for 127 contrasts.

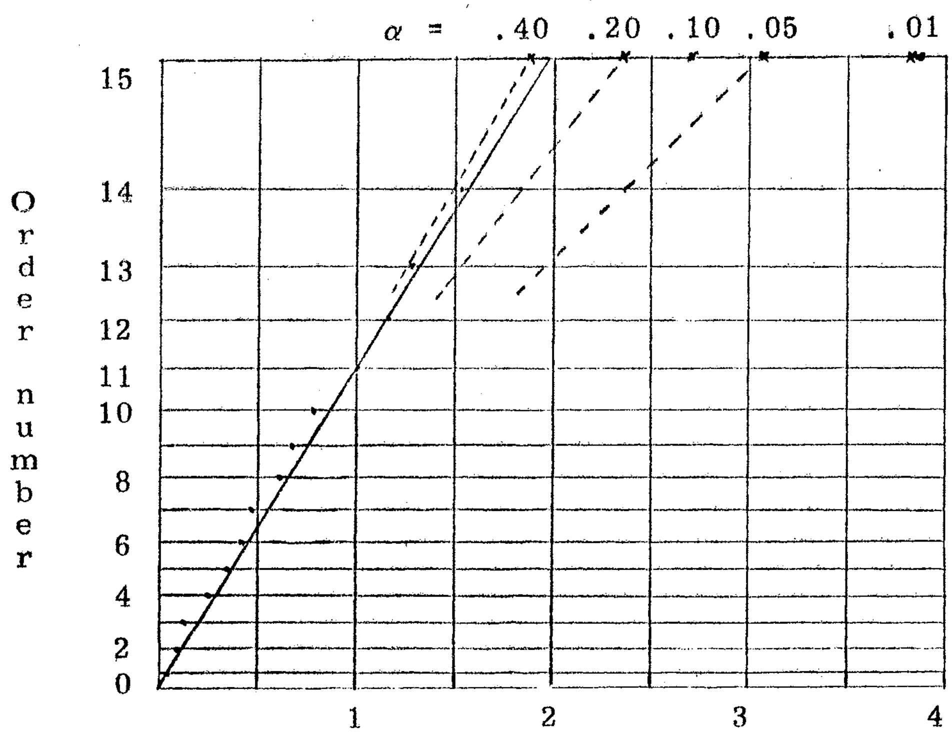

Fig. 15 - An example of a standardized half-normal for 15 contrasts

CHAPTER VI

USE OF HALF-NORMAL PLOTS IN CRITICIZING DATA

It appears that the half-normal plots are sensitive to certain types of defects in experimental data and that they can sometimes be used in detecting these defects.

In an experiment conducted with more than one replication with the usual assumptions that the experimental error is distributed normally with mean μ and variance σ^2, it is expected that the range of the observations on experimental units treated alike will have the half-normal distribution. In other words, the difference between the largest observation and the smallest observation in a block has the half-normal distribution. These ranges can easily be found from the raw data. They can then be ordered and plotted on the appropriate half-normal grid. Any large deviation from the half-normal line would indicate the possible presence of a bad value in the data used to calculate the range of the one that deviated. In short, it would suggest that the largest reading is too large or the smallest reading is too small. It also suggests that the data from this block should not be used in the estimation of the experimental error.

As an example, Davies[6] presents a 2 x 3 x 4 experiment with two replications. There are two methods of manufacturing a product, three temperatures at which it is manufactured, and four batches of raw material from which it is manufactured. Table VIII shows the ordered ranges of the data from the 24 blocks. Under blocks, the first digit indicates the method number, the second the temperature and and the third the batch used. This data is plotted on a special half-normal grid for n equals 24 in Figure 16. Here our σ-estimator is the 17th ordered range or 4. 3. As there are no significant deviations from the half-normal line, it would appear that there are no significantly bad values in the data.

TABLE VIII

Order	Block	Range	Order	Block	Range
24	131	8. 9	12	212	2. 8
23	213	7. 3	11	214	2. 6
22	234	6. 2	10	123	2. 2
21	231	5. 6	9	111	2. 1
20	222	4. 9	8	121	2. 1
19	224	4. 7	7	232	2. 0
18	124	4. 3	6	211	1. 2
17	113	4. 3	5	112	. 8
16	114	4. 2	4	133	. 8
15	223	3. 4	3	221	. 7
14	233	3. 1	2	134	. 6
13	122	3. 0	1	132	. 1

In a 2^p factorial experiment without replication, a single very wide value may be spotted from inspecting the raw data if only one or two main effects are real. However, if there are more real main effects and interactions and the wild value is not too large, then there may be several responses exceeding the bad one. This would make it harder to find. This wild value would appear in every contrast and it would increase half of them by the amount of the bias and the other half would be decreased by the same amount. However, the absolute value of those contrasts near zero would increase and a half-normal plot of the ordered contrasts will have too few contrasts near zero.

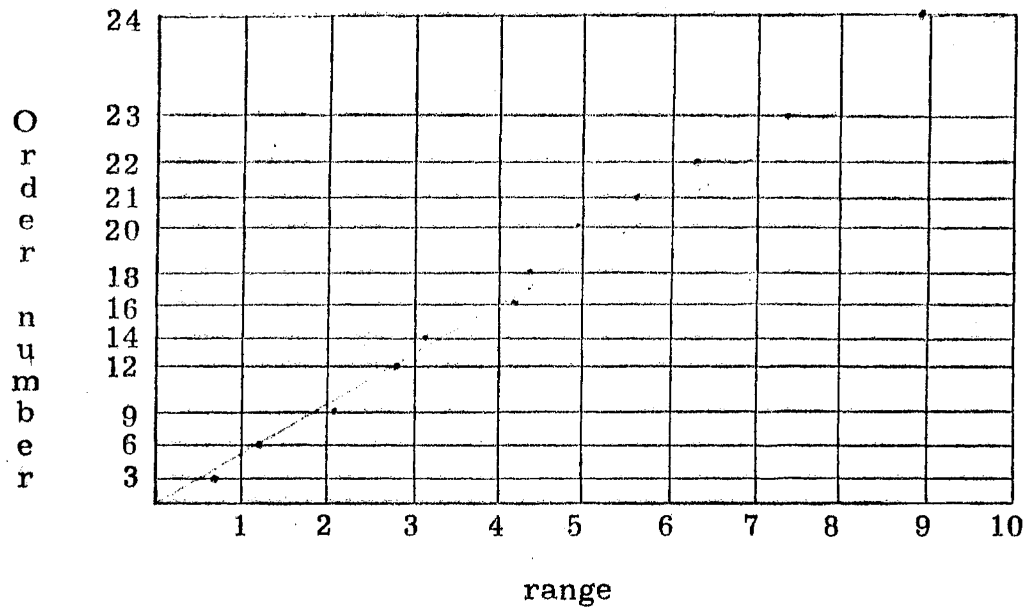

Fig, 16 - Half-normal plot for 24 ranges.

A straight line through the smaller contrasts will not pass through the origin but will pass to the right of it. Figure 17 is the plot of the contrasts from the original data from an experiment with 15 contrasts. As there are few contrasts near zero and a straight line drawn through the 12 smallest would pass to the right of the origin, it was felt that there was a wild value present. From previous knowledge, it was felt that one of the observations was about seven points too large. It was reduced by seven and new contrasts calculated. Figure 18 is the plot of the contrasts from the corrected data. Here, since there is such a wide difference between the 10th and the 11th contrast, and a line through the 10th contrast and the origin seems to be a good half-normal line for the ten smallest contrasts, it is felt that the 10th contrast is a better σ-estimator than the 11th contrast. Therefore, the half-

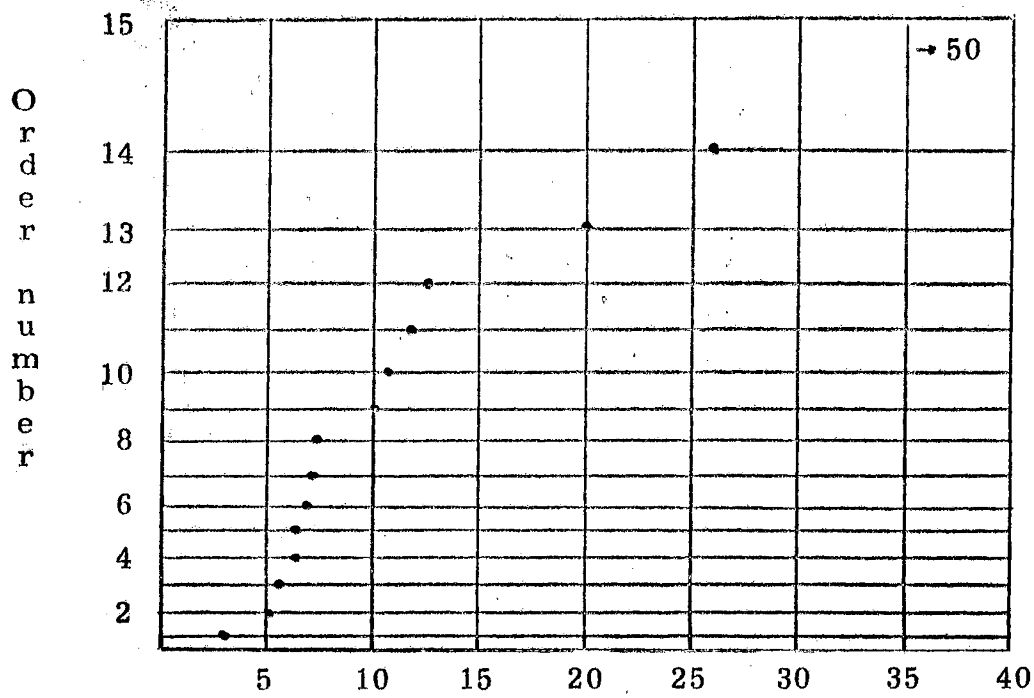

Fig. 17 - Half-normal plot for 15 contrasts from original data.

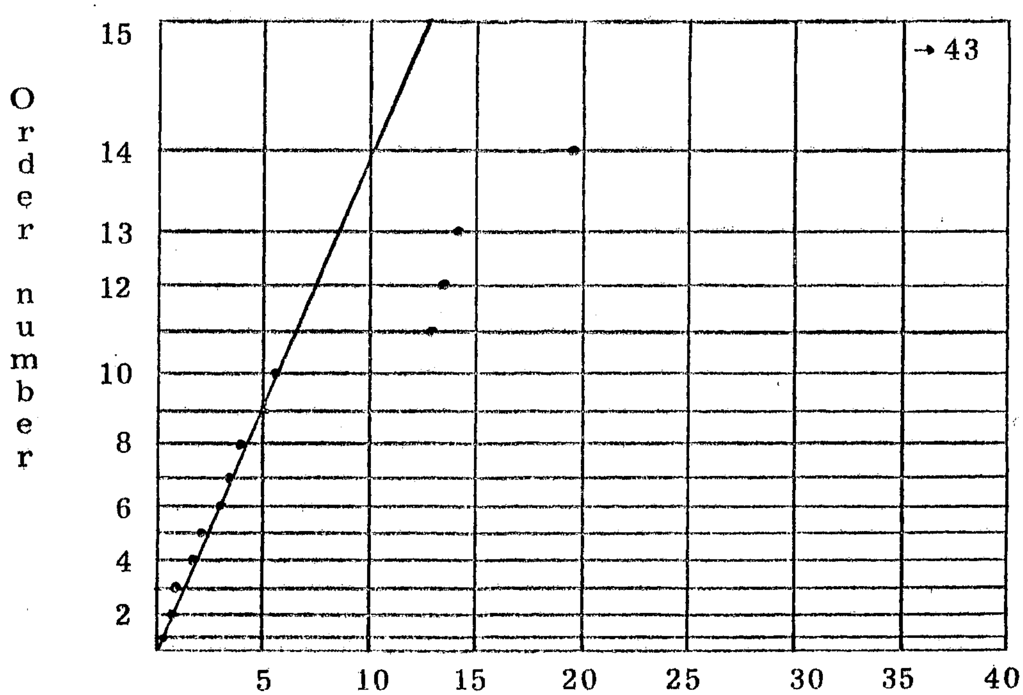

Fig. 18 - Half-normal plot for 15 contrasts from corrected data.

normal line was drawn as shown in Figure 17 and the conclusion reached that the top five effects and interactions are significant.

CHAPTER VII

CONCLUSIONS

While this paper was devoted primarily to the study of the use of the half-normal plots in the criticism and interpretation of 2^p factorial experiments, it seems that it would be equally useful in the interpretation of all types of factorial experiments. While no published study has been made, it seems that it could also be useful in other types of experiments where the assumption of normality has been made.

It must be pointed out, however, that the use of half-normal plots as suggested here is still full of subjective biases. It is not offered as a general substitute for the analysis of variance. Cuthbert Daniel[5] has made half-normal plots for all of the 2^{p-q} factorial experiments in the standard texts and the easily accessible journals as well as several hundred unpublished industrial experiments. In over nine tenths of them, he has found that the lower order contrasts give acceptable linear graphs. However, smoothness of graph is not a guarantee of perfection. A smooth line with no effects securely found might sometimes lead us to over-estimate the magnitude of the error and as a result to miss a number of real effects. An extremely irregular line might lead to the other sort of error, judging effects to be real when in fact they are not.

It appears that there is much work left to be done in the development of the half-normal plot as a useful statistical tool.

BIBLIOGRAPHY

1. Birnbaum, A., On the Analysis of Factorial Experiments without Replication, Technometrics, I, 1959.

2. Chernoff, H., and Lieberman, G. L., Use of Normal Probability Paper, American Statistical Journal, 49, 1954.

3. Daniel, C., Fractional Replication in Industrial Research, Proceedings of Third Berkeley Symposium on Mathematical Statistics and Probability, Vol. V, University of California Press, Berkeley and Los Angeles, 1956.

4. Daniel, C., Locating Outliers in Factorial Experiments, Technometrics, II, 1960.

5. Daniel, C., Use of Half-Normal Plots in Interpreting Factorial Two-Level Experiments, Technometrics, I, 1959.

6. Davies, O. L., Editor: Design and Analysis of Industrial Experiments, Second Edition, Oliver and Boyd, London, and Hafner, New York, 1956.

7. Dixon, W. J., and Massey, F. J., Introduction to Statistical Analysis, Second Edition, McGraw-Hill, 1957.

8. Freund, J. E., Modern Elementary Statistics, Sixth Edition, Prentice Hall, 1952.

9. Goulden, C. H., Methods of Statistical Analysis, Burgess Publishing Company, 1936.

VITA

VADIS GLENN GODBEY

Candidate for the Degree of

Master of Science

Report: THE USE OF HALF-NORMAL PLOTS IN THE INTERPRETATION OF TWO-LEVEL FACTORIAL EXPERIMENTS

Major Field: Mathematics

Biographical:

Personal Data: Born at Bethelridge, Kentucky, July 31, 1928, the son of William Orville and Leora Godbey.

Education: Attended grade school in Argyle and Bethelridge, Kentucky; graduated from Middleburg High School, Middleburg, Kentucky in 1946; received the Bachelor of Arts degree from Berea College, Berea, Kentucky, with a major in Mathematics, in 1951; completed requirements for the Master of Science degree in May, 1963.

Professional Experience: Entered the United States Air Force in 1951, and is now a Captain; greatest part of military experience has been as a Communications Officer at various Air Force establishments in the United States and over seas.