ANALYSIS OF DESIGNED EXPERIMENTS WITH COMPLEX ALIASING

M. Hamada and C.F.J. Wu

IIQP Research Report
RR-91-01

January 1991

ANALYSIS OF DESIGNED EXPERIMENTS WITH COMPLEX ALIASING

M. Hamada and C.F.J. Wu

IIQP Research Report
RR-91-01

January 1991

ANALYSIS OF DESIGNED EXPERIMENTS WITH COMPLEX ALIASING

M. Hamada and C.F.J. Wu

Department of Statistics and Actuarial Science
University of Waterloo
January 1991

ABSTRACT

Traditionally, the Plackett-Burman designs have been used in screening experiments for identifying only important main effects. They have been criticized for their complex aliasing patterns which, according to conventional wisdom, gives confusing results. This paper goes beyond the traditional approach by proposing an analysis strategy that entertains interactions in addition to main effects. The proposed procedure, based on the precepts of effect sparsity and effect heredity, exploits the designs' complex aliasing patterns, thereby turning their "liability" into an advantage. The procedure is demonstrated using data from three real experiments that show the potential for extracting important information available in the data that until now has been missed. Some limitations are discussed and extensions to overcome them are given. The proposed procedure also applies to more general mixed level designs that have become increasingly popular.

Key Words: Interactions, Mixed Level and Plackett-Burman Designs, Partial Aliasing, Screening Experiments.

Introduction

Because of their run size economy, the Plackett-Burman (1946) (PB) designs of 12, 20 and 24 runs are useful experimental plans for studying a large number of two-level factors. For the same reason, mixed level orthogonal arrays such as $L_{18}(2\times3^7)$ and $L_{36}(2^{11}\times3^{12})$ have become increasingly popular for studying both two-level and three-level factors. Except for obvious cases where some interactions are orthogonal to main effects, analysis of such experiments has been confined to main effects only. This is caused by the complex aliasing patterns of these designs. Consider the most "notorious" case, the 12 run PB design with 11 factors; for each factor X, its main effect is partially aliased with the 45 two-factor interactions not involving X, thereby making it difficult to disentangle or interpret the significance of interactions. Most textbooks on design of experiments such as Daniel (1976) recommend PB designs only for screening purposes under the assumption of additivity of the factor main effects. Daniel's reservations are apparent because he refers to the PB designs' complex aliasing patterns as their "hazards" (page 294). Traditional wisdom regarding PB designs can be summarized by the comment made in Draper (1985) that unless (i) the interactions are small or negligible or (ii) there are relatively few "important" factors, the results from PB designs may be confusing.

In practice, PB designs have been used successfully even when the additivity assumption cannot be ascertained. For example, Snee (1985) stated that "it has been my experience that it (referring to the PB designs' complex aliasing patterns) is not a problem and these designs do a good job of identifying those variables that have important effects." In this paper we go beyond the traditional approach by entertaining interactions in PB designs (and other designs with complex aliasing) that can be identified and estimated with reasonable precision. Our rationale is based on the empirical evidence that in many experimental investigations, the variation exhibited in the data is largely attributable to only a few effects. This phenomenon may be called the Pareto Principle in experimental design or effect sparsity as Box and Meyer (1985) refer to it. Since only few main effects and even fewer two-factor interactions are relatively important, the complex aliasing patterns of these designs can be simplified greatly, thereby making it possible to study a few interactions.

In contrast with designs having complex aliasing patterns, the 2^{k-p} fractional factorial designs have straightforward aliasing patterns; i.e., any two effects are either orthogonal or completely aliased. While the number of aliased effects with any main effect remains small, there is also a disadvantage. For example, in the cast fatigue experiment studied later, the FG interaction and D main effect are identified as important by an experiment using a 12 run PB design. Suppose that an eight run 2^{7-4} resolution III design had been used instead with D=FG as one of its defining relations. Then, it would have been impossible to disentangle FG from its alias D. On the other hand, both D and FG are estimable from the PB design. Moreover, this problem may not disappear by doubling the run size from 8 to 16. The heat exchanger experiment also studied later provides such an example. If a 16 run 2^{10-6} resolution III design had been used to study the ten factors with EG=H as one of its defining relations, then it would have been impossible to disentangle H and EG. In contrast, the 12 run PB design that was actually used allows E, H, EG and EH to be estimated simultaneously. Thus, this new perspective of PB designs and more general designs with complex aliasing allows their known deficiency to be turned into an advantage.

After presenting an analysis of an experiment on weld repaired castings, we introduce and discuss the notions of *effect sparsity* and *effect heredity.* Based on them, we propose an analysis strategy to study effectively main effects and interactions for general designs with complex aliasing when there are only a few significant interactions that are partially aliased with the main effects. Reanalyses of three actual experiments presented in Examples 1 to 3 demonstrate the tremendous gain in information that can be achieved by using the proposed procedure. For example, identifying one or two interactions in the three experiments increases the R^2 by 98%, 89% and 45%, respectively.

The proposed procedure can miss some important main effects when there are several significant interactions that are partially aliased with the main effects. Two extensions are proposed for overcoming this problem. Both the problems and remedies are illustrated in Examples 4 and 5.

The paper concludes with a discussion of the procedure's applicability to other designs with complex aliasing; these designs include larger run PB designs and the mixed level designs studied in Wang and Wu (1990, 1991).

Finally, we comment on Taguchi's view regarding the analysis of designs with complex aliasing. In Taguchi and Wu (1979, page 35), he states that " ... no interactions are calculated even if they exist. ... these interactions are treated as errors, so it is advantageous to have the effects of these interactions uniformly distributed in all (design matrix) columns." From this and other statements made elsewhere, he seems to believe that estimated main effects are not affected by interactions because they are smeared or evenly spread across all the design matrix columns. Our analyses of some real experiments and hypothetical examples show that this is generally untrue. In particular, ignoring interactions can lead to (i) important effects being missed, (ii) spurious effects being detected and (iii) estimated effects having reversed signs resulting in incorrectly recommended factor levels.

Example 1: Weld Repaired Cast Fatigue Experiment

Hunter, Hodi and Eager (1982) used a 12 run PB design to study the effects of seven factors on fatigue life of weld repaired castings. The factors and levels are listed in Table 1. (The seven factors were assigned using the first seven columns of the design matrix given in Table 2 with "+" and "-" denoting the high and low factor levels.) In the original analysis, heat treat (factor D) and polish (factor F) were identified as significant, although heat treat had a much smaller effect (with a p-value around .2); the R^2 for the model (D,F), which consists of the D and F main effects, is .59. Our initial analysis using a half-normal probability plot (Figure 1) suggested that only polish is significant (R^2=.45), however. Then, by entertaining all the interactions with polish, we found a significant polish by final treat (FG) interaction. Adding FG to F doubles the R^2 to .89! Moreover, adding the heat treat effect (D) identified by Hunter et al. (1982) to the model (F, FG) only increases the R^2 to .92; i.e., the heat treat factor appears to be insignificant. A summary of these analyses given in Table 3 shows the substantial improvement by including the interaction effect. The improved analysis may also lead to a better understanding of the process and some benefits as discussed below. Using the model (F,FG), we obtain predicted fatigue lifetimes by

$$\hat{y} = 5.7 + .46F - .46FG, \tag{1}$$

where F and G take 1 or -1 corresponding to the "+" and "-" levels, respectively. By changing F from chemical ("-") to mechanical ("+") with G set at peening ("+"), the predicted life is increased by 19% (=.92/4.78). The additional increase in predicted life by changing G from peening ("+") to none ("-") is 16% (=.92/5.70). In practice, this prediction has to be verified by a confirmation experiment. Nevertheless, the potential for a dramatic improvement would not have been possible without discovering the FG interaction. An additional benefit comes from simplifying the original process by elimination of the peening step. Next, we present a general analysis strategy for designs with complex aliasing patterns. Later, we return to this example to discuss additional reasons for supporting model (1).

Table 1: Factors and Levels for Cast Fatigue Experiment

	factor	level +	level -
A.	initial structure	β treat	as received
B.	bead size	large	small
C.	pressure treat	HIP	none
D.	heat treat	solution treat/age	anneal
E.	cooling rate	rapid	slow
F.	polish	mechanical	chemical
G.	final treat	peen	none

Table 2: Design and Data for Cast Fatigue Experiment

	design											logged
run	A	B	C	D	E	F	G	8	9	10	11	data
1	+	+	-	+	+	+	-	-	-	+	-	6.058
2	+	-	+	+	+	-	-	-	+	-	+	4.733
3	-	+	+	+	-	-	-	+	-	+	+	4.625
4	+	+	+	-	-	-	+	-	+	+	-	5.899
5	+	+	-	-	-	+	-	+	+	-	+	7.000
6	+	-	-	-	+	-	+	+	-	+	+	5.752
7	-	-	-	+	-	+	+	-	+	+	+	5.682
8	-	-	+	-	+	+	-	+	+	+	-	6.607
9	-	+	-	+	+	-	+	+	+	-	-	5.818
10	+	-	+	+	-	+	+	+	-	-	-	5.917
11	-	+	+	-	+	+	+	-	-	-	+	5.863
12	-	-	-	-	-	-	-	-	-	-	-	4.809

Table 3: Model Summary for Cast Fatigue Experiment

model	R^2
F	.45
F, D	.59
F, FG	.89
F, FG, D	.92

The Proposed Analysis Strategy

Our analysis strategy is based on two precepts. First, we assume that only a small number of factorial effects are important relative to the rest. This can be called the Pareto Principle in experimental design or *effect sparsity* as Box and Meyer (1985) refer to it and has been empirically verified in many experiments. Second, we assume that when a two factor interaction is significant, at least one of the corresponding factor main effects is also significant. We refer to this precept as *effect heredity* (a term suggested by R. Sitter). It is often difficult to provide a good physical interpretation for a significant AB interaction without A or B being significant.

The two precepts form the basis of our analysis strategy which is described as follows:

Step 1 Entertain all the main effects and interactions that are orthogonal to the main effects. Use standard analysis methods such as analysis of variance or half-normal plots to select significant effects. Go to Step 2.

Step 2 Using effect heredity, entertain (i) the effects identified in the previous step and (ii) the two-factor interactions that have at least one component factor appearing in the effects in (i). Also, consider (iii) interactions suggested by the experimenter. Use a forward selection regression procedure to identify significant effects among the effects in (i)-(iii). Go to Step 3.

Step 3 Use a forward selection regression procedure to identify significant effects among the effects identified in the previous step as well as all the main effects. Go to Step 2.

Iterate between Steps 2 and 3 until the selected model stops changing. (Effect sparsity suggests that only a few iterations will be required.)

Step 4 Assess the model by looking at residual diagnostics.

Note that the standard analysis of PB designs ends at Step 1. If all two factor interactions are entertained indiscriminately in Step 2, it is possible to get a good fitting model with only interaction terms and no main effects; hence, nonsensical models may be obtained without assuming effect heredity.

Some explanation for handling factors with more than two levels is needed. Here, we discuss only three level factors. For quantitative factors, the linear and quadratic effects can be entertained to assess the effects of linearity and curvature on the response. The corresponding regression covariates are formed by assigning the following values to the three levels, $(1/\sqrt{2})(-1, 0, 1)$ for the linear effect and $(1/\sqrt{6})(1, -2, 1)$ for the quadratic effect. For qualitative factors, the same covariates can be used and interpreted as a comparison between the first and third levels for the linear effect and a comparison between the average of the first and third levels with the second level for the quadratic effect. More flexibility is gained by comparing other levels such as the first and second level by using the values $(1/\sqrt{2})(-1,1,0)$ for the "linear" effect. Since interactions between two three-level factors require four degrees of freedom, initially entertaining the linear by linear component (as is done in response surface modeling) frees up three degrees of freedom so that other effects can be considered.

If a more exhaustive search of models is desired, the final model obtained by the proposed procedure can be viewed as a good starting model. Other models can be entertained by adding or deleting effects one at a time. For example, while (A,B,BC) is chosen by forward selection, (A,B,AB) with a slightly smaller R^2 may be preferred if it renders a better interpretation.

Results from Real Experiments

Next, we use the proposed procedure to reanalyze data from three actual experiments. These examples demonstrate the opportunity for identifying additional important effects which the traditional approach misses.

Cast Fatigue Experiment Revisited

Hunter et al. (1982) noted a discrepancy between their findings and previous work, namely, the sign of the heat treat (factor D) effect was reversed. They suggested that the cause was an interaction between heat treat and cooling rate (DE) and claimed that the design did not generate enough information to determine this. Because one property of this design is that a factor main effect is orthogonal to all interactions involving the factor, we know that the DE interaction does not affect the sign of the factor D main effect. The design's aliasing patterns with some additional calculations help to explain this apparent reversal, however. From Table 4 which gives estimates of the 11 main effects from Step 1 and their partial aliases with FG, we observe that $\hat{D}$ (= -.516) actually estimates $D + \frac{1}{3}FG$. Consequently, we first need to correct the D estimate as follows: $\hat{D} - \frac{1}{3}\hat{FG} =$

$(-.516-(-.310))$ = -.206. While the corrected estimate remains negative, a 95% confidence interval, (-.522,.110)=($-.206 \pm .316$), shows that a positive D main effect is possible.

Note from Table 4 that the A, B and C estimated effects are insignificant and are more likely explained by an FG interaction since $\pm\frac{1}{3}\hat{FG} = \mp.31$ is close to their estimates, .33, .29 and -.25, respectively. Also, the effects A-E and 8-11 have the same sign as their FG aliases, which lends further support to the existence of a significant FG interaction.

Interestingly, Hunter et al. (1982) suggested an FG interaction to explain a discrepancy between the estimated effect for factor G and other results obtained from optical micrographs. Our statistical analysis provides quantitative support of this engineering conjecture; namely, model (1) suggests that peening increases the lifetime of chemically polished casts but reduces the lifetime of mechanically polished ones.

Table 4: Effects and Aliases for Cast Fatigue Experiment

effect	estimated effect	alias pattern
A	.326	$A - \frac{1}{3}FG$
B	.294	$B - \frac{1}{3}FG$
C	-.246	$C + \frac{1}{3}FG$
D	-.516	$D + \frac{1}{3}FG$
E	.150	$E - \frac{1}{3}FG$
F	.915	F
G	.183	G
8	.446	$-\frac{1}{3}FG$
9	.453	$-\frac{1}{3}FG$
10	.081	$-\frac{1}{3}FG$
11	-.242	$+\frac{1}{3}FG$

Example 2: Blood Glucose Experiment

Henkin (1986) used an 18 run mixed level design to study the effect of eight factors on blood glucose readings of a clinical laboratory testing device as shown in Table 5. Here we consider only one aspect of the study which was to identify factors that affect the mean reading, y.

Table 5: Design and Data for Blood Glucose Experiment

				design					mean
run	A	G	B	C	D	E	F	H	reading
1	1	1	1	1	1	1	1	1	97.94
2	1	1	2	2	2	2	2	2	83.40
3	1	1	3	3	3	3	3	3	95.88
4	1	2	1	1	2	2	3	3	88.86
5	1	2	2	2	3	3	1	1	106.58
6	1	2	3	3	1	1	2	2	89.57
7	1	3	1	2	1	3	2	3	91.98
8	1	3	2	3	2	1	3	1	98.41
9	1	3	3	1	3	2	1	2	87.56
10	2	1	1	3	3	2	2	1	88.11
11	2	1	2	1	1	3	3	2	83.81
12	2	1	3	2	2	1	1	3	98.27
13	2	2	1	2	3	1	3	2	115.52
14	2	2	2	3	1	2	1	3	94.89
15	2	2	3	1	2	3	2	1	94.70
16	2	3	1	3	2	3	1	2	121.62
17	2	3	2	1	3	1	2	3	93.86
18	2	3	3	2	1	2	3	1	96.10

For each three-level factor, we decomposed its main effect into linear and quadratic effects. Since the interaction between the two-level factor A and the three-level factor G is orthogonal to all other main effects, we entertained 17 effects in Step 1 consisting of A, the linear and quadratic effects of B to H, and the linear and quadratic effects of AG. Figure 2 presents a half-normal plot of the 17 estimated effects which identifies the quadratic effects of factors E and F as significant (R^2=.36). In Step 2, we expanded the candidate variables by entertaining all the interactions involving E or F. A forward selection procedure identified the EF linear by linear interaction as significant. Adding this interaction to the main effects model found in Step 1 substantially increased the R^2 to .68, yielding the fitted model

$$\hat{y} = 95.9 + 18.3E_q - 17.0F_q - 23.7E_l F_l. \tag{2}$$

Since factors E and F are used in calibrating the machine, identifying the EF linear by linear interaction may be important for setting up the machine to give unbiased readings.

Example 3: Heat Exchanger Experiment

Specht (1985) used a 12 run PB design to study the importance of 10 factors $(A-H, J, K)$ on heat exchanger lifetime, where lifetime is defined as the number of cycles until the unit develops a tube wall crack. The data are actually interval censored data arising from eight inspections. For this paper, we used the interval midpoints as the response, which had no effect on the conclusions. See Hamada and Wu (1991) for a method to formally handle censored data. The data and the design matrix can be found in Specht (1985) and Hamada and Wu (1991).

In Step 1, a half-normal plot identified only the factor E main effect as important (R^2=.62). In Step 2, the nine interactions between E and the other factors were entertained using a forward selection procedure with EG and EH being identified as important; the (E,EG,EH) model has an R^2=.90, a gain of .28! In fact, a more extensive search entertaining all 10 main effects and 45 two-factor interactions gave the same model, yielding a predicted lifetime of

$$\hat{y} = 70.9 - 21.5E + 11.1EG - 9.2EH\ . \tag{3}$$

A table similar to Table 4 with the alias structure of the 11 estimated effects in terms of the 10 main effects, EG and EH appeared in Hamada and Wu (1991). It can be used to explain why the estimates for B, K, C and the unassigned column were larger than those of the remaining factors in the half-notmal plot from Step 1.

In Specht's analysis, only the main effect E was detected. Nevertheless, the data contained information about two important interactions which the proposed procedure found. Using model (3), the levels (E low, G low, H high) are recommended for increasing the lifetime. One of the experimental runs had this combination of E, G and H and indeed had the largest lifetime among the experimental runs! This agreement with the experimental result lends further credibility to model (3). Thus, arbitrarily setting levels for G and H by ignoring important interactions could lead to an inferior design with a shorter lifetime.

Limitations and Some Extensions

The proposed method works well when only a few of the interactions that are partially aliased with the main effects are significant and smaller than the main effects. Otherwise some significant main effects may be missed as the next two examples show. In Example 4, two interactions are twice as large as the main effect A. In Example 5, both the interactions BC and CD are partially aliased with the main effect A. Since we know the true model in both these examples, we can determine how well the proposed method identifies the real effects in the model. After the examples, we discuss extensions of the proposed method for overcoming this problem.

Example 4. Assume that the 12 run PB design given in Table 2 is used to study a process where the true model is $y = A + 2AB + 2AC + \epsilon$, (A,B,C) $= \pm 1$ according to the design and the other main effects B-K are zero. A normally distributed error ϵ with mean 0 and standard deviation .25 was used to obtain the data. From Figure 3, the proposed procedure identifies E, H, I and K as significant in Step 1. Clearly, because A, B and C are not detected, none of the real effects will be detected in the procedure's remaining steps.

Example 5. Assume that a 12 run PB design is used to study a process where the true model is $y = 2A + 4C + 2BC - 2CD + \epsilon$. A normally distributed error ϵ with mean 0 and standard deviation .5 was used to obtain the data. From Figure 4, the proposed procedure identifies only C in Step 1, whose corresponding R^2 is .700. In Step 2, by entertaining two factor interactions between C and the other factors, BC is identified as important; the R^2 of the model (C,BC) is .764. In Step 3, when the other main effects (A,B,D-K) are entertained again, we find that the three models (C,G), (C,H) and (C,J) have nearly the same R^2's as (C,BC). Moreover, by adding an additional effect, we find that the R^2's of (C,BC,CD), (C,G,H), (C,H,J) and (C,G,J) are nearly the same with values of .877, .863, .844 and .855, respectively. Thus, the proposed procedure may find several incompatible models that are equally plausible.

The A main effect is missed by the half-normal plot in Step 1 because A is smaller than the two interactions in Example 4 and is masked by the two interactions of equal size in Example 5. One solution is to relax the criterion for significance in Step 1 so that more main effects are included. It then becomes less likely to miss the real main effects because their estimates, though masked by their partial aliases, are usually among the larger ones in the main effect analysis. Relaxing the criterion can also result in the detection of spurious effects in Step 1, however. For example, suppose that the criterion is relaxed so that the five largest effects are identified (including A) in Step 1 for Example 4. Then, in Step 2, the correct model (A,AB,AC) is identified when (AB,AC) are entered into the model while (I,K,E,H) are dropped.

Another solution to this problem is to replace Step 1 by the following step, which we call Step 1′:

Step 1′ For each factor X, entertain X and all its interactions XY with other factors. Use a model selection criterion such as C_p (Mallows 1973) to find the best model from these candidate variables denoted by M_X. Repeat this for each of the factors and then choose the best model from all the M_X's. Go to Step 3.

Then iterate between Steps 3 and 2 as in the proposed procedure.

The second solution is motivated by Example 4 and indeed identifies the true model in one step; i.e., Step 1′ identifies the three effects A, AB and AC.

Although these two extensions enable the proposed procedure to cover a broader class of models, there is still a limitation to the information provided by the data and design. If the procedure finds several incompatible but equally plausible models such as those found in Example 5, it is a strong indication that no analysis method can distinguish between them. One way to resolve the ambiguity is to conduct additional runs with the goal of discriminating between these models. Strategies for generating additional runs to de-alias effects or to discriminate between competing models can be found in Daniel (1976, chapter 14) and Box, Hunter and Hunter (1978, page 413). Alternatively, one can use a D-optimal design algorithm such as DETMAX (Mitchell 1974) or ACED (Welch 1984). The algorithm can search for runs that, together with the existing runs, are D-optimal for a model consisting of all effects of interest including those to be de-aliased.

Applications to Other Designs

Although the proposed procedure has been demonstrated primarily with the 12 run PB design, it can also be used for other PB designs such as the 20, 24 and 28 run designs. Their aliasing patterns are qualitatively similar to those for the 12 run design. For example, each main effect in the 20 run design is partially aliased with 153 two factor interactions; 144 have coefficient $\pm.2$ and 9 have -.6. Similarly, the procedures can be used for mixed level designs as was demonstrated for the blood glucose experiment which used an L18 design. The aliasing pattern of the L18 design with one two-level factor and seven three-level factors is simpler than the 12 run PB design, however. For example, out of the 21 linear by linear interactions between any two three-level factors, only six are partially aliased with the two-level factor main effect (with coefficient -.5). Since the percentage of partially aliased effects is smaller for the L18 design, problems such as missed real effects, detected spurious effects and model discrimination ambiguity are less serious. The proposed methods can also be used for general mixed level orthogonal designs such as those considered in Wang and Wu (1991).

To achieve further run savings, Wang and Wu (1990) considered designs whose main effects are either orthogonal or nearly orthogonal. Since the estimated main effects can be correlated, standard analysis methods in Step 1 such as analysis of variance or half-normal plots need to be replaced by forward selection or best subset selection procedures. Except for this change, the rest of the procedure (including Step 1′) remains the same.

Acknowledgements

We thank R. J. MacKay for helpful comments on an earlier version. This research was supported by General Motors of Canada Limited, the Manufacturing Research Corporation of Ontario, and the Natural Sciences and Engineering Research Council of Canada.

References

Box, G.E.P., Hunter, W.G. and Hunter, J.S. (1978). *Statistics for Experimenters*, John Wiley and Sons, New York, NY.

Box, G. and Meyer, R. D. (1985). "Some New Ideas in the Analysis of Screening Designs." *Journal of Research of the National Bureau of Standards* 90, pp. 495-502.

Daniel, C. (1976). *Applications of Statistics to Industrial Experiments*, John Wiley and Sons, New York, NY.

Draper, N.R. (1985). "Plackett and Burman Designs" in *Encyclopedia of Statistical Sciences*, John Wiley and Sons, New York, NY, pp. 754-758.

Hamada, M. and Wu, C.F.J. (1991). " Analysis of Censored Data from Highly Fractionated Experiments." *Technometrics*, to appear in February.

Henkin, E. (1986). "The Reduction of Variability of Blood Glucose Levels" in *Fourth Supplier Symposium on Taguchi Methods*, American Supplier Institute, Inc., Dearborn, MI, pp. 758-785.

Hunter, G.B., Hodi, F.S. and Eager, T.W. (1982). "High-Cycle Fatigue of Weld Repaired Cast Ti-6Al-4V." *Metallurgical Transactions* 13A, pp. 1589-1594.

Mallows, C.L. (1973). "Some Comments on C_p." *Technometrics* 15, pp. 661-675.

Mitchell, T.J. (1974). "An Algorithm for the Construction of D-Optimal Experimental Designs." *Technometrics* 16, pp. 203-210.

Plackett, R.L. and Burman, J.P. (1946). " The Design of Optimum Multifactorial Experiments." *Biometrika* 33, pp. 305-325.

Snee, R.D. (1985). "Experimenting with a Large Number of Variables" in *Experiments in Industry*, American Society for Quality Control, Milwaukee, WI, pp. 25-35.

Specht, W. (1985). "Heat Exchanger Product Design Via Taguchi Methods" in *Third Supplier Symposium on Taguchi Methods*, American Supplier Institute, Inc., Dearborn, MI, pp. 302-318.

Taguchi, G. and Wu, Y. (1979). *Introduction to Off-Line Quality Control*, Central Japan Quality Control Association, Nagoya, Japan.

Wang, J.C. and Wu, C.F.J. (1990). "Nearly Orthogonal Arrays with Mixed Levels and Small Runs." University of Waterloo Institute for Improvement in Quality and Productivity Research Report RR-90-03.

Wang, J.C. and Wu, C.F.J. (1991). "An Approach to the Construction of Asymmetrical Orthogonal Arrays." *Journal of the American Statistical Association*, to appear.

Welch, W.J. (1984). "Computer-Aided Design of Experiments for Response Estimation." *Technometrics* 26, pp. 217-224.

Figure 1. Half-Normal Plot of Step 1
for Cast Fatigue Experiment

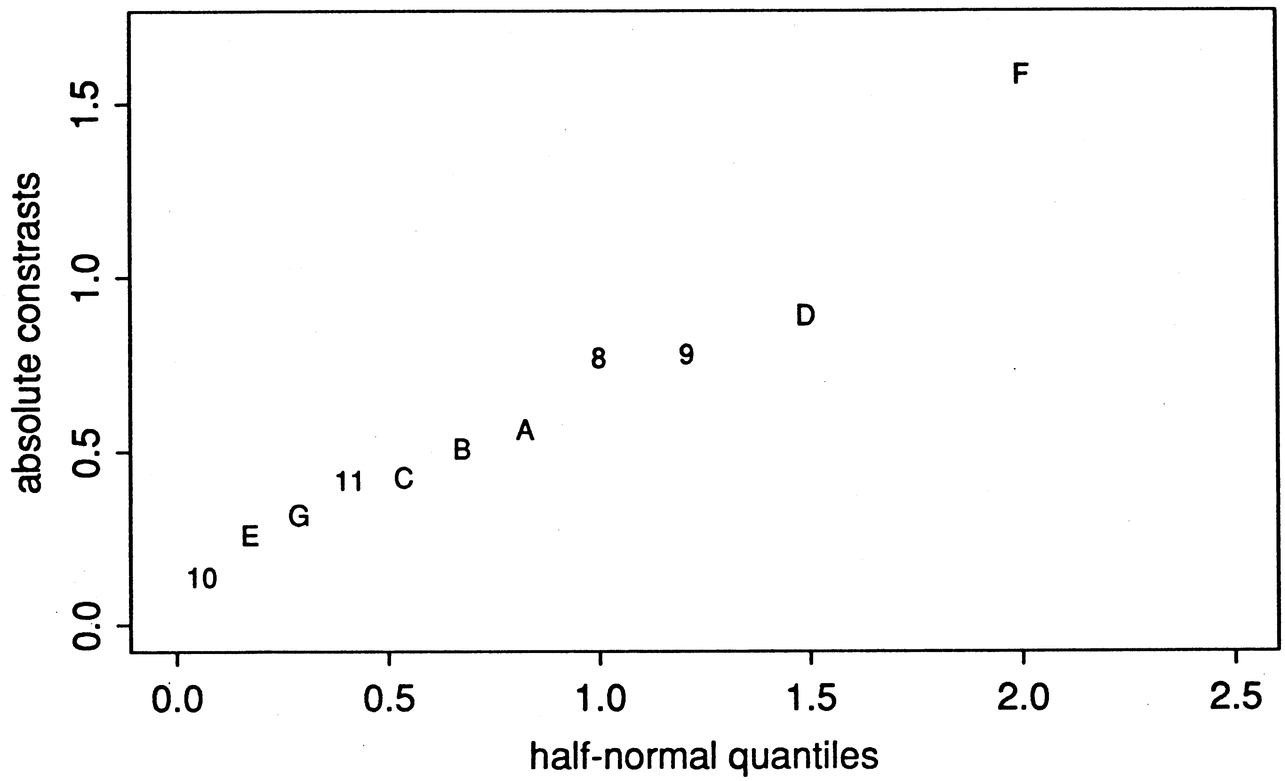

Figure 2. Half-Normal Plot of Step 1
for Blood Glucose Experiment

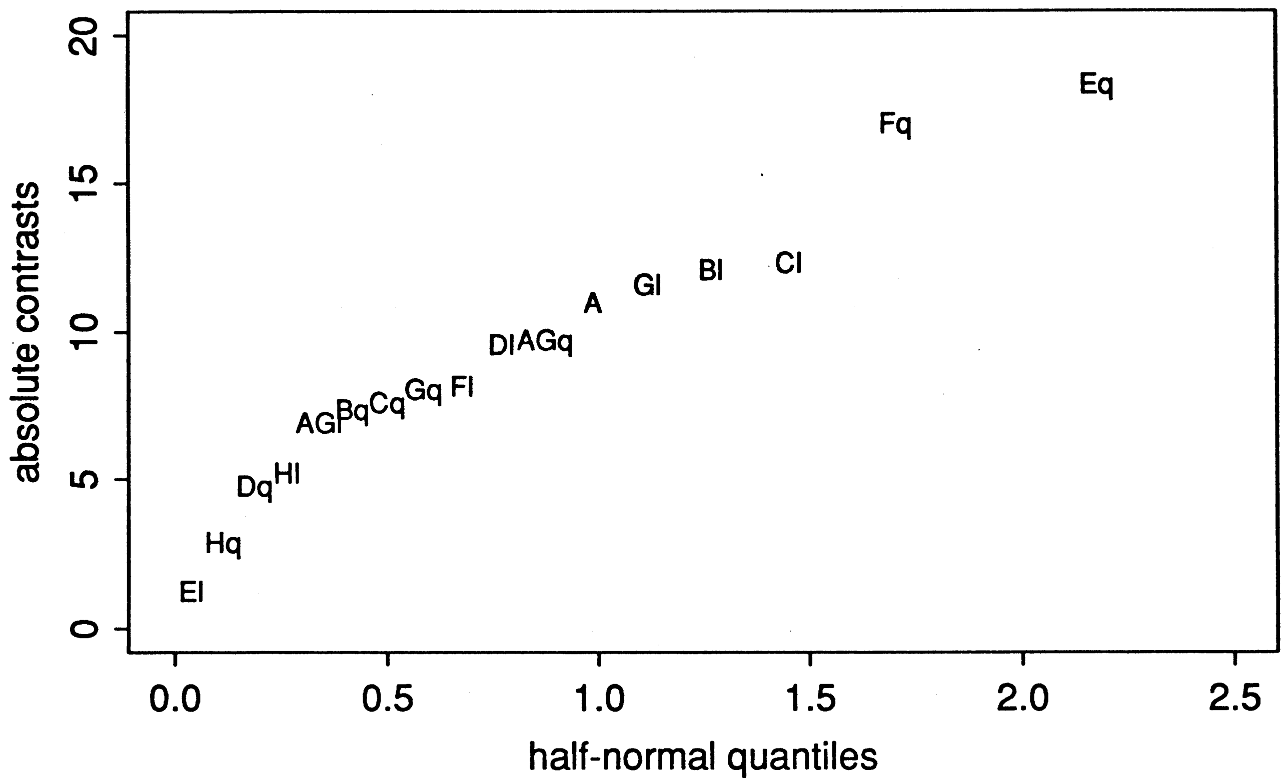

Figure 3. Half-Normal Plot for Step 1
from Example 4, Y=A+2AB+2AC

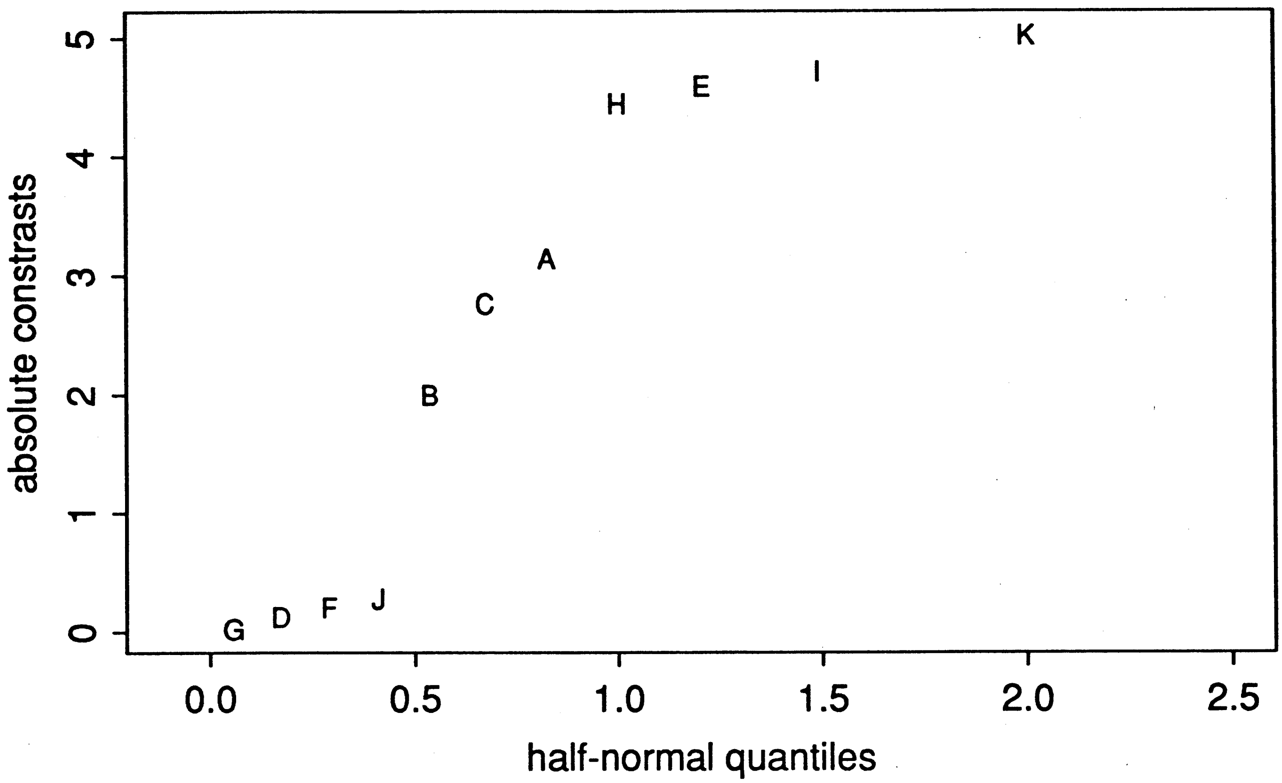

Figure 4. Half-Normal Plot for Step 1
from Example 5, Y=2A+4C+2BC-2CD