Statistica Sinica 2(1992), 313-334

ROBUST PARAMETER DESIGN WITH UNCONTROLLED NOISE VARIABLES

Anne E. Freeny and Vijayan N. Nair

AT&T Bell Laboratories, Murray Hill

Abstract: Robust parameter design, originally proposed by Taguchi, is a very useful tool for reducing variation and improving product/process quality. By exploiting the relationships between control parameters (design factors) and noise variables, it reduces the effect of uncontrollable variations on the response. This is done by using statistically designed experiments in off-line situations where the settings of the noise variables are controlled and systematically introduced and their relationships with design factors studied.

Frequently, however, it is too expensive or not practically feasible to control the noise variables, even in off-line experiments. Variations in the noise variables can then invalidate the usual methods of analysis. In this paper, we develop alternative methods of analysis for situations where the noise variables are uncontrolled but can be observed. Our approach involves treating the noise variables as covariates and modeling both the location parameters and the regression coefficients as functions of the design factors. These coefficients can be viewed as the equivalent of Taguchi's S/N ratios for reducing variability induced by the observed noise variables. We propose a general data-analysis strategy for determining various dispersion and location effects and improving performance under this framework. The approach is also applicable to experiments where there are covariates, and one must remove the effects of these nuisance variables before identifying the location and dispersion effects of the design factors. The ideas are illustrated by applying them to an experiment for thermal design of cabinets for telecommunications switching systems.

Key words and phrases: Design of experiments, dispersion effects, location effects, quality improvement, S/N ratios.

1. The Problem

Let Y denote a quality characteristic associated with a product or process of interest. Following Taguchi (1986, 1987), we can express

$$Y = f(\boldsymbol{x}; \boldsymbol{n}) \tag{1.1}$$

where $\boldsymbol{x}$ denotes control/design factors, i.e., factors that are easily controlled by the process/product engineer and are fixed during manufacturing and/or operation, $\boldsymbol{n}$ denotes variables that are not easily controlled and may vary during manufacturing and/or operation, and $f(\cdot)$ denotes the functional relationship. We shall call $\boldsymbol{x}$ and $\boldsymbol{n}$ "design factors" and "noise variables" respectively. Variations in the noise variables $\boldsymbol{n}$ lead to variations in Y, the measure of product/process quality. Traditional approaches to controlling or reducing this variability, such as acceptance sampling and tolerance design, are costly and do not lead to any real improvements in quality. The novel idea behind Taguchi's robust parameter design is to first try to reduce the effect of the variations in the noise variables on the response, instead of using costly measures to actually control variations in the noise variables. This is done by exploiting the functional relationship $f(\cdot)$ to determine the setting of the design factors, not only to get the average performance on target but also to minimize variability around this target value.

For robust parameter design to be successful, the functional relationship $f(\cdot)$ in (1.1) must be such that the design factors and noise variables "interact". For if an underlying additive model

$$g(Y) = f_1(\boldsymbol{x}) + f_2(\boldsymbol{n}) \tag{1.2}$$

holds, the settings of $\boldsymbol{x}$ have no influence on variability.

In practice $f(\cdot)$ is either unknown or, as in the case of computer experiments, very complex. Thus one conducts statistically designed experiments, either physical or computer experiments, and obtains the responses at the experimental settings of the design factors. For robust parameter design experiments, the noise variables are also usually controlled, typically in off-line environments, and their effects studied systematically. More specifically, in these experiments

$$Y_{ij} = f(\boldsymbol{x}_i; \boldsymbol{n}_j, \varepsilon_{ij}),\ i = 1, \ldots, m,\ j = 1, \ldots, k, \tag{1.3}$$

where Y_{ij}'s are the responses, $\boldsymbol{x}_i$ is the ith row of the design matrix (corresponding to the design factors), the identified noise variables $\boldsymbol{n}$ are controlled at fixed levels and the settings $\boldsymbol{n}_j$, $j = 1, \ldots, k$, are repeated for each of the m design settings, and ε_{ij} denote the remaining, unidentified noise variables, including measurement error.

The usual analyses of location and dispersion effects are based on

$$\bar{Y}_i = k^{-1} \sum_{j=1}^{k} Y_{ij} \tag{1.4}$$

and

$$S_i^2 = k^{-1} \sum_{j=1}^{k} (Y_{ij} - \bar{Y}_i)^2 \tag{1.5}$$

respectively. Taguchi actually recommends, instead of (1.5), a different measure (S/N ratio) for the "on-target" problem. His measure implicitly assumes a multiplicative model where the variance is proportional to the square of the mean (see Léon et al. (1987)). In such situations, we prefer the use of the approximately equivalent procedure of applying a logarithmic transformation to the data and using (1.5) based on the transformed data. See Nair and Pregibon (1986) and Box (1988) for more details. Throughout, we assume that when variance-stabilizing transformations are necessary, they have been carried out and that we are dealing with the transformed data.

The analyses above depend critically on the fact that the same values of the noise variables are repeated across the m design settings. In many situations, however, some of the important noise variables cannot be carefully controlled. This is certainly true for on-line experiments. Even in off-line situations, it may be too expensive or not practically feasible to control the noise variables. For example, in studying performances of computer or communication networks, the primary noise variables are system and network load conditions. (See Phadke (1989) for an application to tuning the performance of a computer system.) These load conditions cannot be controlled in most experiments. Variations in the noise variables can then invalidate the conclusions obtained from the usual analyses based on (1.4) and (1.5). If the mean levels of the noise variables vary from one design setting to another, they can confound estimates of location effects based on (1.4). Similarly, differences in their variances will confound the estimates of dispersion effects based on (1.5). These problems will be demonstrated explicitly in the next few sections when we consider specific models and data.

In this paper, we consider situations where the noise variables can vary but are observable. (In the computer/communication network applications, for example, it is possible to measure the various load conditions and so adjust for the changes in the noise variables.) We propose a general data analysis strategy which is based on modeling the responses directly. Our approach involves treating the noise variables as covariates and modeling both the location parameters and the (regression) coefficients as functions of the design factors. This allows us to determine the interactions between the design factors and the observed noise variables and to exploit them to reduce the effect of this source of variability. Variability due to unobserved noise variables can be identified by analyzing the squared residuals from the fitted model. The approach presented here is also applicable to experimental situations with covariates, where one has to remove the effects of these nuisance variables before identifying the important location and dispersion effects.

The paper is organized as follows. In Section 2, we use a real application on thermal design of cabinets for telecommunications switching equipment to motivate the problem and the issues. Section 3 develops the underlying concepts and models for a single observed noise variable. The proposed data analysis strategy is outlined in Section 4 and is illustrated by applying it to the thermal design experiment. Section 5 deals with several generalizations, including the case of multiple noise variables.

The direct modeling of responses as a function of design factors and noise variables has also been considered by Welch et al. (1990), Shoemaker et al. (1991), and Lucas (1990). Our approach simplifies to the formulations discussed by these authors when the noise variables are controlled and the models are all linear. In general, however, there are differences in modeling, analysis, interpretations, and applications. The methods developed here remove the need to strictly control the noise variables, and hence one can do robust parameter design experiments more easily in on-line environments, provided the noise variables can be observed. Observing and measuring noise variables during on-line production and operation have other benefits and can result in better insights on the causes of variation.

2. An Illustrative Example

An experiment was conducted at AT&T Bell Laboratories to study the surface temperature within cabinets in telecommunications switching equipment. Recent technological advances have considerably increased the power dissipated by circuit packs in these switching equipment, well beyond original guidelines. The goal of the experiment was to determine if circuit packs can be kept within their temperature limits by designing the cabinets appropriately with suitable numbers of fans and shelves and by careful placement of high power dissipation circuits. We shall use the context of this real experiment but with synthetic data to motivate and illustrate the problem and issues.

A 2^4 full factorial experiment was conducted with four factors and four replications. The four factors were: A – power dissipation on the shelf of interest; B – cumulative power dissipation on all the shelves below the shelf of interest; C – number of fans; and D – the shelf of interest. There are six shelves in a cabinet, numbered from the lowest to the highest. The fans are located below shelf 1. Table 1 shows the levels of the four factors.

Table 1. Factors in thermal design experiment

Factor Code	Factor	Levels
A	Power Dissipation (watts per circuit pack)	0 – 40
B	Cumulative Power Dissipation (watts per circuit pack)	0 – 200
C	Number of Fans	1 – 3
D	Shelf Location	2 – 6

Surface temperatures were measured at three locations in the cabinet – backplane, center, and faceplate. We shall restrict attention here to just surface temperatures measured at the center. Table 2 gives the (synthetic) surface temperature data used in our analyses. The temperatures (in celsius) have been centered by subtracting the overall mean of the ambient temperatures. The order of the observations corresponds to the order of the standard 2^4 factorial design matrix with the factors $A - D$ in Yates' order.

Table 2. Surface temperature

	Replicate					
Run	1	2	3	4	$\bar{Y}_i$	S_i^2
1	−0.1	−0.3	−1.0	−1.0	−0.57	0.22
2	15.8	14.4	16.7	16.0	15.75	0.93
3	2.1	3.1	0.8	2.9	2.25	1.09
4	23.0	24.1	23.3	24.5	23.75	0.48
5	0.0	0.4	1.2	1.4	0.78	0.44
6	13.9	13.4	13.1	12.8	13.33	0.22
7	9.4	8.6	10.2	9.0	9.33	0.47
8	27.1	25.9	26.6	27.1	26.70	0.32
9	−1.1	−1.5	0.7	0.4	−0.35	1.18
10	17.8	18.0	17.2	18.5	17.90	0.29
11	15.0	11.9	21.0	18.9	16.73	16.42
12	23.7	22.0	29.8	26.0	25.40	11.39
13	1.1	2.2	2.5	3.2	2.28	0.76
14	16.7	16.2	18.3	17.7	17.25	0.90
15	9.3	6.9	9.8	6.7	8.20	2.57
16	24.0	24.8	28.1	27.3	26.08	3.84

The usual approach to identify important dispersion effects is to fit a log-linear model to S_i^2 in (1.5) as a function of the design factors. Figure 1 is a half-normal plot of the effects computed from this analysis. It suggests that D, the number of shelves, and to a lesser extent B, the cumulative amount of power dissipated, are important factors in controlling variability in surface temperature.

In reality, however, the four observations at each design setting were not true replications. The surface temperature within the cabinet is affected by the ambient temperature, and it was not possible to control the ambient temperature at fixed levels. Instead their values were recorded. Table 3 gives the *actual* ambient temperatures that were observed during this experiment, centered by subtracting their overall mean. We have reordered the observed values slightly for illustrative purposes. We shall regard the data in Table 3 as the observed values of the noise variable.

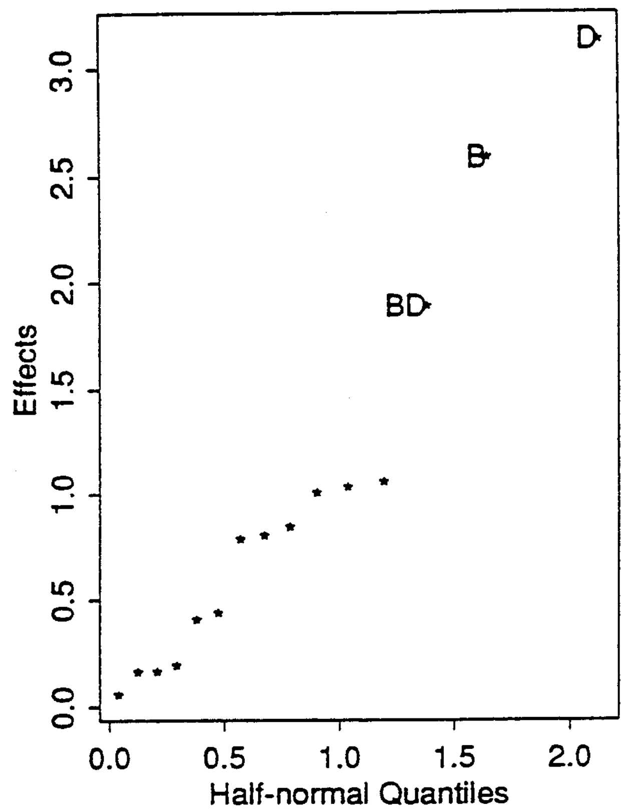

Figure 1. Half-normal plot of effects of log (variance) of surface temperature. The 15 effects are the result of an analysis of variance fitting the full factorial model $\log(S_i^2) = A * B * C * D$.

Table 3 shows that $S_i^2(n)$, the "within-run" variances of the ambient temperatures, range over two orders of magnitude. To understand how these differences could have potentially confounded the analysis in Figure 2, we analyzed the $S_i^2(n)$'s as a log-linear function of the design factors. Figure 2 shows the half-normal plot of the effects obtained from this analysis, and it suggests that the effect of factor D is strongly confounded with the differences in $S_i^2(n)$. Hence the conclusion from Figure 1 that D has a large dispersion effect could be misleading.

One quick-and-easy way to try to account for the differences in $S_i^2(n)$ is to analyze, instead of (1.5), the ratio $S_i^2/S_i^2(n)$. (In the next section, we will consider a model under which the use of this measure can be justified.) Figure 3 is the half-normal plot of the effects computed by fitting a log-linear model to this ratio. We see that the conclusions from this analysis are different from those from Figure 2. D is no longer the most important effect, and we might even conclude that there are no important effects. We shall see in the next two sections, however, that the conclusions from this analysis can also be misleading, and that B, C, and BC are important dispersion effects for this data set. Another disadvantage of this approach is that it does not extend easily to multiple noise variables and other more complex situations.

Table 3. Ambient temperature

	Replicate					
Run	1	2	3	4	$\bar{n}_i$	$S_i^2(n)$
1	−0.57	−0.77	−0.67	−0.67	−0.67	0.01
2	−0.47	−0.57	0.33	−0.37	−0.27	0.17
3	−2.47	−1.87	−2.57	−2.27	−2.30	0.10
4	−0.97	−0.77	−0.87	−0.67	−0.82	0.02
5	0.23	0.23	1.13	0.83	0.60	0.20
6	−2.67	−2.87	−3.27	−2.97	−2.95	0.06
7	−0.57	−0.87	0.03	−0.77	−0.55	0.16
8	0.83	0.53	0.53	1.23	0.78	0.11
9	−0.77	−1.07	0.93	1.03	0.03	1.22
10	2.63	2.03	1.73	2.73	2.28	0.23
11	1.43	0.53	3.43	3.03	2.10	1.85
12	−0.77	−1.17	1.23	−0.17	−0.22	1.10
13	2.13	2.23	2.73	3.43	2.63	0.35
14	0.73	0.33	2.13	1.33	1.13	0.61
15	−0.57	−2.67	−0.67	−2.87	−1.70	1.55
16	−0.77	−0.87	1.63	−0.17	−0.05	1.34

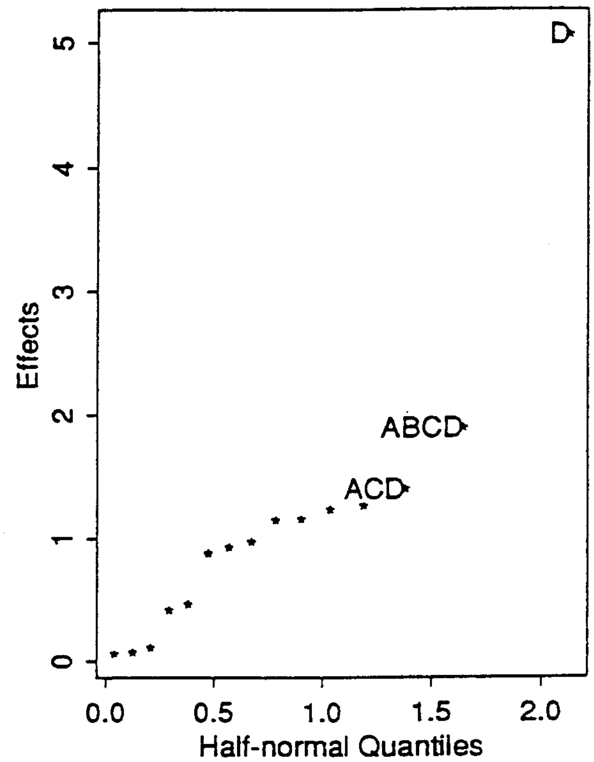

Figure 2. Half-normal plot of effects of log (variance) of ambient temperature. The 15 effects are the result of an analysis of variance fitting the model $\log(S_i^2(n)) = A*B*C*D$.

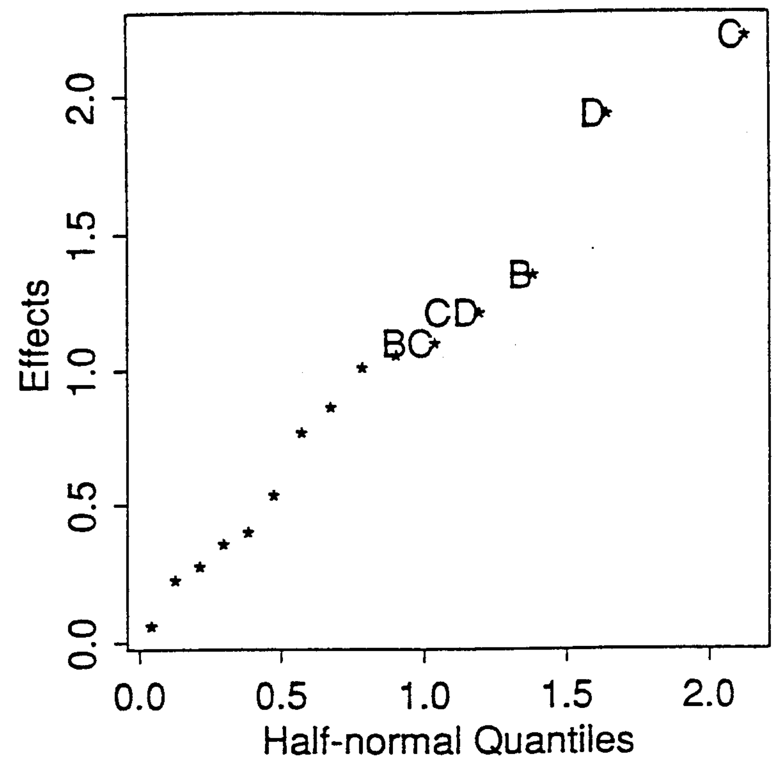

Figure 3. Half-normal plot of effects of log (ratio of variances). The 15 effects are the result of an analysis of variance fitting the model $\log(S_i^2/S_i^2(n)) = A * B * C * D$.

What alternative measure to consider will depend on the situation of interest. One extreme case, (a), is where ambient temperature is the primary noise variable, i.e., most of the variation in surface temperature is caused by variations in ambient temperature. The focus of the analysis should then be on modeling the relationship between surface and ambient temperatures and choosing the settings of the design factors to minimize the effect of this relationship. The other extreme case, (b), is where ambient temperature is only a covariate in the experiment and is not a noise variable in practice. This will be the case if, during operation, the switching equipment is located in air-conditioned buildings and the ambient temperature does not vary much. So, the ambient temperature in the experiment is just a nuisance variable whose effect on the observed surface temperatures must be removed. We can then analyze the variances of the residuals from the fitted model, instead of (1.5), to determine the dispersion effects. The intermediate case, (c), is where ambient temperature is just one of the important noise variables and there may be other, unidentified noise variables whose effects we should also understand. In this case, we should do both (a) and (b). In subsequent sections, we consider these issues in more detail and propose a general data analysis strategy.

We can examine in a similar fashion how the differences in the mean levels of the ambient temperatures affected the analysis of location effects based on (1.4). It turns out, however, that the differences in mean ambient temperatures were small in comparison to the differences in surface temperatures, and did not influence the analysis significantly in this case.

3. Models for a Single Noise Variable

3.1. Constant variance

Suppose that, after a suitable transformation of the response if necessary, we have the model

$$Y_{ij} = \mu_i + g(n_{ij};\beta_i) + \varepsilon_{ij}, \tag{3.1}$$

for $i = 1,\ldots,m$ and $j = 1,\ldots,k$, where $g(\cdot)$ is a fixed function. Further, for suitable link functions $h_0(\cdot)$ and $h_1(\cdot)$,

$$h_0(\mu_i) = \boldsymbol{x}_i'\boldsymbol{\alpha}, \tag{3.2}$$

$$h_1(\beta_i) = \boldsymbol{x}_i'\boldsymbol{\phi}, \tag{3.3}$$

and

$$\mathrm{Var}(\varepsilon_{ij}) \equiv \sigma^2. \tag{3.4}$$

Here "i" corresponds to the design setting $\boldsymbol{x}_i$ (i.e., ith row of the design matrix) and "j" corresponds to the "replications" within the ith setting, and n_{ij}'s are the observed values of the noise variable. We assume that there are at least two replications, i.e., $k \geq 2$.

The constant variance assumption in (3.4) will be relaxed later. Note, however, that this assumption implicitly implies that other unidentified noise variables do not interact with design factors, so that we cannot choose the settings of the design factors to reduce the effect of these noise variables on the response. This assumption may be reasonable if most of the important noise variables are identified and studied explicitly during the experiment.

We shall assume that the function $g(\cdot)$ in (3.1) is centered at the mean of the noise distribution, i.e., $g(0;\beta_i) \equiv 0$. So, $\mu_i = \mu(\boldsymbol{x}_i)$ measures the mean of the response distribution at the design setting $\boldsymbol{x}_i$. The effect of changes in the noise variable on response is measured by $\beta_i = \beta(\boldsymbol{x}_i)$, together with $g(\cdot)$. We should, therefore, choose the design factor setting (and hence β_i) appropriately to minimize the effect of potential changes in this noise variable during manufacturing/operation. Thus, β_i is the appropriate measure for assessing variability in this situation, where the noise variable is uncontrolled and varies from one design setting to another. We can thus view it as the analog of Taguchi's "SN-ratio" for this problem.

Consider the special case of a linear regression in (3.1), that is

$$Y_{ij} = \mu_i + \beta_i n_{ij} + \varepsilon_{ij}. \tag{3.5}$$

If the ε_{ij}'s are small,

$$\beta_i^2 \approx S_i^2 / S_i^2(n), \tag{3.6}$$

where S_i^2 is the "within-run" variance of Y_{ij} given by (1.5), and $S_i^2(n)$ is the "within-run" variance of the noise variable. In the special case where the n_{ij}'s are fixed and the same settings, $n_{ij} = n_j$, $j = 1, \ldots, k$, are repeated across the m design settings, $S_i^2(n) \equiv S^2(n) =$ constant. So, analyzing β_i is equivalent to analyzing S_i^2, the usual measure in (1.5). When the noise variables vary, however, the two are not equivalent. But in this case, the relationship (3.6) suggests that one could use instead the ratio $S_i^2/S_i^2(n)$. This provides a justification for the analysis in Figure 2 in the last section. The problem with this ratio, however, is that it actually estimates $\beta_i^2 + \sigma^2/S_i^2(n)$. So unless σ^2 is small, the measure can still be affected by the variations in $S_i^2(n)$.

Going back to the general set-up in (3.1), equations (3.2) and (3.3) specify the relationships between the design factors and the μ_i's and β_i's respectively. Let $\boldsymbol{\alpha} = (\alpha_0, \alpha_1, \ldots, \alpha_{p-1})$, where p is the rank of the design matrix, and let the ϕ_j's be defined similarly. We call α_j, $j \geq 1$, the location effects and ϕ_j, $j \geq 1$, the dispersion effects of the design factors. More precisely, the ϕ_j's are dispersion effects associated with the noise variable n_{ij}. The magnitudes of these ϕ_j's measure the deviation of (3.3) from the null hypothesis $H_0 : \beta_i \equiv$ constant (ϕ_0), i.e., there is no interaction between the design factors and the noise variable. Thus, ϕ_j, $j \geq 1$, represent the (generalized) interaction between the design factors and the noise variable in this setting. When the models in (3.1)–(3.3) are all linear and the noise variable is controlled at n_j, $j = 1, \ldots, k$, these parameters specialize to the usual interaction effects, and our approach then reduces to that discussed in Welch et al. (1990), Shoemaker et al. (1991), and Lucas (1990).

3.2. Unequal variances

So far, we have assumed that there are no important dispersion effects that can be attributed to other, unobserved noise variables. To allow for this more general case, we let

$$\text{Var}(\varepsilon_{ij}) = \sigma_i^2, \; i = 1, \ldots, m \tag{3.7}$$

where

$$\nu(\sigma_i) = \boldsymbol{x}_i'\gamma. \tag{3.8}$$

So, for $j \geq 1$, γ_j's measure the deviation from the null hypothesis of constant variance $(\sigma_i \equiv \gamma_0)$, and indicate which of the factors are influential in reducing variability caused by unobserved noise variables. In practice, it may be difficult to estimate these effects unless we have pure replications or effect sparsity, i.e., only a small proportion of the effects are active.

There is one important situation where the primary dispersion effects of interest will be γ. This is the case where the n_{ij}'s are covariates in the experiment but do not vary during manufacturing or are otherwise do not affect the variation in practice. In our thermal design application, for example, if the equipment is located in air-conditioned bulidings, ambient temperature does not vary much during operation. The ϕ_j's in (3.3) are then not relevant for reducing variation; they are only important for estimating the β_i's and removing the effects of ambient temperatures on the observed surface temperatures in the experiment.

4. Data Analysis Strategy

We first outline the proposed strategy and then illustrate it by applying it to the thermal design experiment. The strategy relies extensively on the use of data-analytic techniques for model identification. These are discussed in the final sub-section. Some of these recommendations are tentative in nature, and further work is needed to refine them.

4.1. The steps

1. Determine the functional form of $g(\cdot)$ in (3.1).

2A. Estimate the β_i's and μ_i's in (3.1) separately for $i = 1, \ldots, m$.

2B. Treating the $\hat{\beta}_i$'s and $\hat{\mu}_i$'s as responses, fit the models (3.2) and (3.3) respectively and obtain preliminary estimates of the location and dispersion effects.

2C. Use graphical methods to identify the active location and dispersion effects, and hence the appropriate submodels in (3.2) and (3.3).

3. Substitute these identified submodels for location and dispersion effects in the combined model

$$Y_{ij} = h_0^{-1}(\boldsymbol{x}_i'\boldsymbol{\alpha}) + g(n_{ij}; h_1^{-1}(\boldsymbol{x}_i'\boldsymbol{\phi})) + \varepsilon_{ij}, \tag{4.1}$$

re-estimate the parameters using (non-linear) least-squares, and compute fitted values $\hat{Y}_{ij}$ and standard errors.

4. To identify any dispersion effects associated with unobserved noise variables in (3.8), compute the studentized residuals r_{ij} from the fitted model in the previous step, and let $\hat{\sigma}_i^2 = \sum_{j=1}^{k} r_{ij}^2/k$. Treating $\hat{\sigma}_i$ as the response, fit the model in (3.8) as a function of the design factors and identify the important effects.

5. If the previous step reveals that the variances are unequal, refit the model in Step 3 using weighted least-squares to allow for these unequal variances, and compute the new fitted values and standard errors. (It may be necessary to iterate this process.)
6. Use the conclusions of the data analysis to determine improved settings of design factors and to predict the improvements in performance.

4.2. Illustration of the strategy

We illustrate the strategy by reanalyzing the data in Section 2.

Step 1. Figure 4 is the plot of the surface temperatures versus the ambient temperatures, Y_{ij} vs n_{ij}, $j = 1, \ldots, 4$ for each of the 16 design settings. The separate least squares lines have been superimposed.

This figure suggests that the linear model in (3.5) is reasonable for settings where the range of ambient temperatures is relatively large. For other settings, the slopes of the least squares lines are not well determined as they have large variances. The functional relationship for these settings cannot be determined precisely in any case. So we will conclude that the model in (3.5) is adequate here. We also see from Figure 4 that the estimated slopes vary over the design settings, suggesting that the β_i's are not constant and that there are dispersion effects.

Step 2A. In general, the parameters must be estimated by nonlinear least-squares. Here, because the assumed model is linear, we can fit least-squares lines to the data at the individual design settings and estimate the parameters as

$$\hat{\beta}_i = k^{-1} \sum_{j=1}^{k} \{[Y_{ij} - \bar{Y}_i][n_{ij} - \bar{n}_i]\} / S_i^2(n) \tag{4.2}$$

and

$$\hat{\mu}_i = \bar{Y}_i - \hat{\beta}_i \bar{n}_i. \tag{4.3}$$

Step 2B. We assume here that $h_0(\cdot)$ and $h_1(\cdot)$ are known and are the identity functions. (See, however, the discussion in the next sub-section.) Let $\boldsymbol{X}$ be the design matrix corresponding to a 2^4 full factorial design and let $\boldsymbol{x}_i$ be its ith row.

To get at the dispersion effects, we fit the model

$$\hat{\beta}_i = \boldsymbol{x}_i' \boldsymbol{\phi} + \nu_i, \tag{4.4}$$

where $\boldsymbol{x}_i'\boldsymbol{\phi}$ is the expected value of $\hat{\beta}_i$ and ν_i is the random component. The variance of $\hat{\beta}_i$ is inversely proportional to $S_i^2(n)$ in Table 3. Since these range over two orders of magnitude, we should take these differences into account by fitting (4.4) using weighted least squares. Therefore, if $\boldsymbol{D}$ is a diagonal matrix with entries $\boldsymbol{D}_{ii} = S_i^2(n)$,

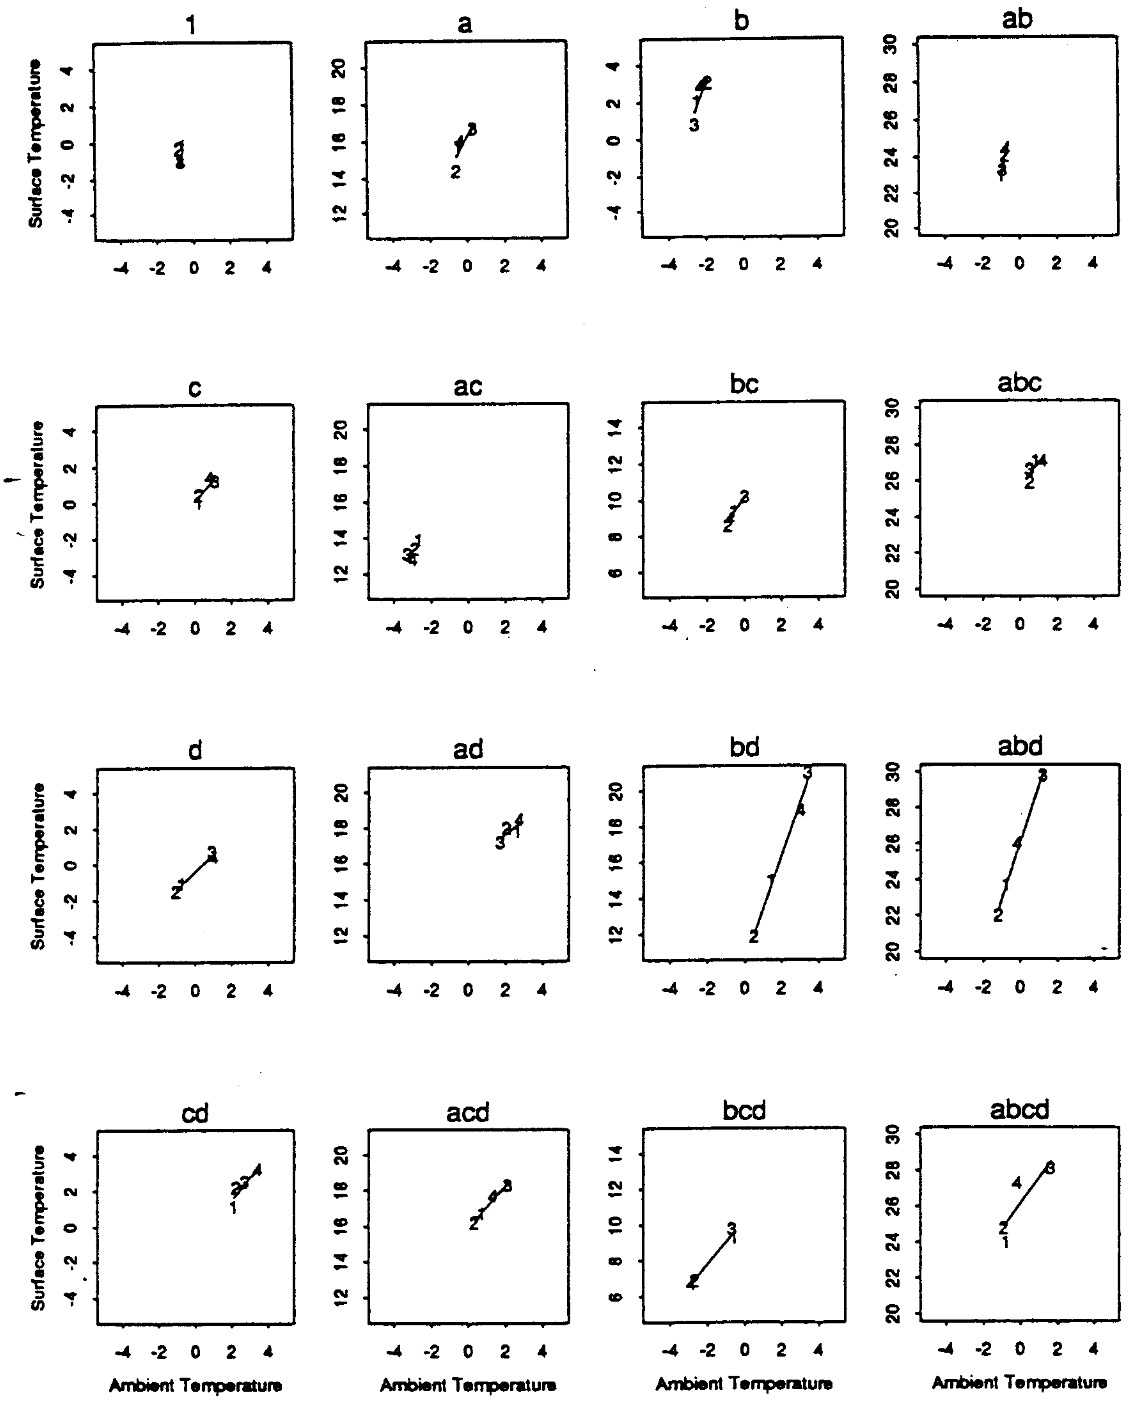

Figure 4. Surface temperature vs ambient temperature for the 16 design settings. Identification above each subfigure gives levels of the factors. The presence of the letter denotes the high level of the factor, its absence denotes the low level. Thus subfigure ab shows the data for the setting with high levels of factors A and B and low levels of factors C and D. Within each subfigure the four replicates are coded 1 to 4. The least squares line for the four points in each subfigure is shown. The y-axis is scaled to have the same range of 10 units for all subfigures.

$$\hat{\phi} = (X'D^{-1}X)^{-1}(X'D^{-1}\hat{\beta}) \tag{4.5}$$

with variance-covariance matrix

$$\mathrm{COV}(\hat{\phi}) = \sigma^2(X'D^{-1}X)^{-1}. \tag{4.6}$$

In the preliminary (model-identification) stage, one would usually start off with a fully saturated model in (4.4), i.e., the design matrix X is $m \times m$ and of full rank. In this case, the weighted least squares estimators in (4.5) are the same as the usual ordinary least squares estimators, but their variances are still given by (4.6).

For identifying location effects, we fit the model

$$\hat{\mu}_i = x_i'\alpha + w_i, \tag{4.7}$$

where $x_i'\alpha$ is the expected value of $\hat{\mu}_i$ and w_i is the random component. For this data set, the differences between the $\hat{\mu}_i$'s and $\bar{Y}_i$'s are small. Since the latter are uncorrelated and have equal variance (and are also uncorrelated with $\hat{\beta}_i$'s), we will use ordinary least-squares to estimate the parameters in (4.7). In general, however, one would have to use generalized least-squares based on the estimated variance-covariance matrix.

Step 2C. We see from (4.6) that the estimators of dispersion effects are correlated, so the usual half-normal plot should not be used for model selection. We use instead a C_p-plot based on weighted least-squares analysis (see Mallows (1973) for use and interpretation). Figure 5 is the C_p-plot for our data.

A careful examination of the plot suggests that the four-term model – constant term plus B, C and BC – is the most reasonable. Note also that the BC interaction term appears throughout and is the single most important effect. So, even if the B and C main effects were not statistically significant, we should still fit the four-term model above so that the results are interpretable.

The conclusions from Figure 5 are different from those based on Figures 1 and 3. It is clear that factor D was identified as important in Figure 1 due entirely to differences in $S_i^2(n)$. Even Figure 3, which tries to account for these differences, leads to wrong conclusions in this case. So we do not recommend the quick-and-easy analysis based on the ratio $S_i^2/S_i^2(n)$.

For identifying location effects, we can use a half-normal plot since, for our data, the $\hat{\mu}_i$'s are approximately uncorrelated and have equal variance. If the situation is different, one would have to use an alternative model selection technique at this stage also. Figure 6 is the half-normal plot of the effects computed from fitting the model (4.7) by ordinary least squares.
Figure 6 suggests that A and B are the two important location effects.

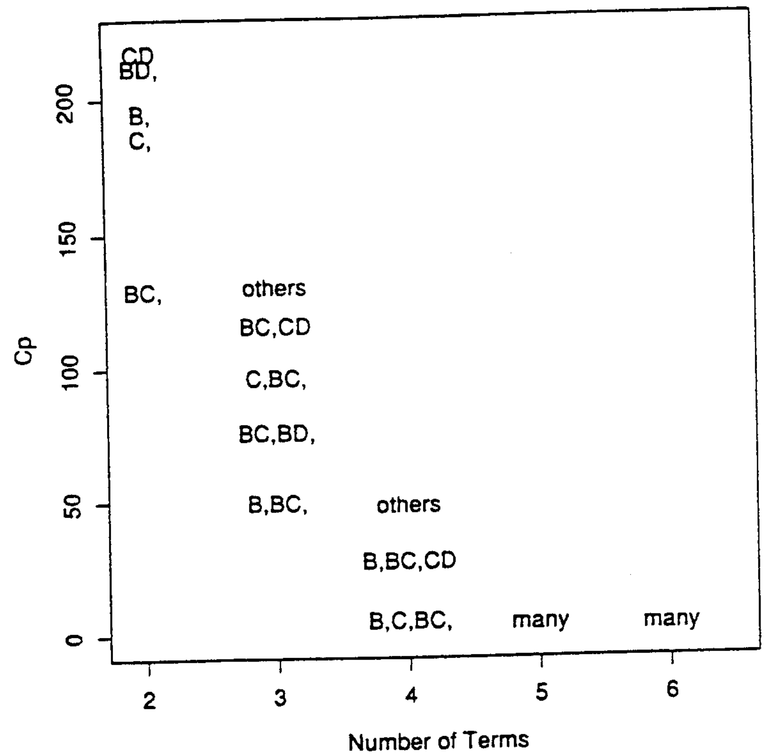

Figure 5. C_p-plot of weighted fits of $\hat{\beta}_i$ to $\boldsymbol{X}$. The number of terms includes the intercept term. Weights are $1/S_i^2(n)$. The plot indicates that BC is the best fitting single term, B and BC, the best fitting two terms, and B, C, and BC the best fitting three terms; and that no further addition of terms will improve the fit.

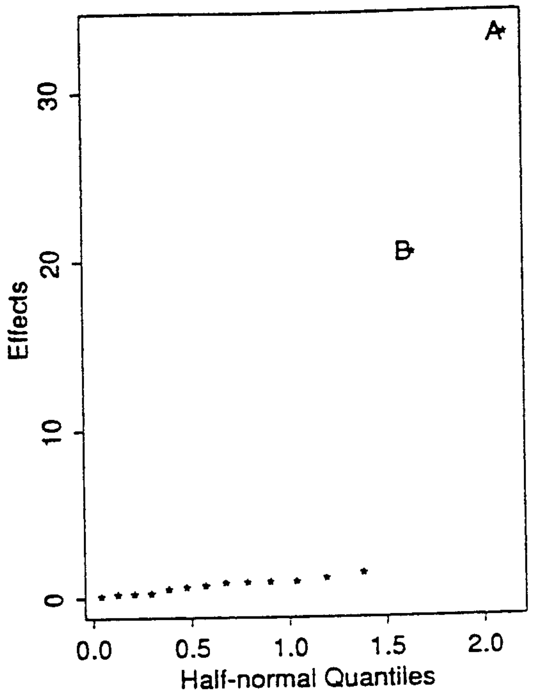

Figure 6. Half-normal plot of location effects. The 15 effects are the result of an analysis of variance fitting the model $\hat{\mu}_i = A * B * C * D$.

Step 3. We substitute the location and dispersion submodels in the overall model in (4.1). Since all the functions are linear, we just have the combined linear model

$$Y_{ij} = \alpha_0 + \alpha_A + \alpha_B + n_{ij}(\phi_0 + \phi_B + \phi_C + \phi_{BC}) + \varepsilon_{ij}. \tag{4.8}$$

Table 4 gives the estimated coefficients and their standard errors obtained by re-estimating the parameters of this sub-model.

Table 4. Estimated coefficients for model (4.8)

Coefficient	Estimate	Standard Error
α_0	13.00	0.062
α_A	8.10	0.070
α_B	5.01	0.062
ϕ_0	1.54	0.042
ϕ_B	0.64	0.042
ϕ_C	−0.50	0.040
ϕ_{BC}	−0.55	0.046

Step 4. The analyses thus far have been based on the assumption of constant variance in (3.4). To see if the more general model in (3.8) holds, we compute the studentized residuals $r_{ij} = (Y_{ij} - \hat{Y}_{ij})/(1 - \nu_{ij})$, where $\hat{Y}_{ij}$'s are the fitted values, and ν_{ij}'s are the diagonal elements of the "hat matrix" for the combined regression problem in (4.8). From this we obtain $\hat{\sigma}_i^2 = \sum_{j=1}^{k} r_{ij}^2/k$. Assuming a log-linear model in (3.8) – see next section for a discussion – we compute the effects associated with the design factors. Figure 7 is a half-normal plot of the effects obtained from this, and it suggests that there are no important dispersion effects associated with unobserved noise variables.

Step 5. Since the assumption of constant variance appears reasonable, we do not have to refit the model at this stage.

Step 6. One should integrate the conclusions from the statistical analyses with subject matter expertise to arrive at suitable factor levels to improve performance. We shall consider, here, just results of the statistical analyses. It should also be reiterated that we are dealing with synthetic data here, so the discussions below are meant to be only suggestive of how one proceeds at this stage.

Table 5 summarizes the effects of the active factors on the $\hat{\beta}_i$'s and $\hat{\mu}_i$'s.

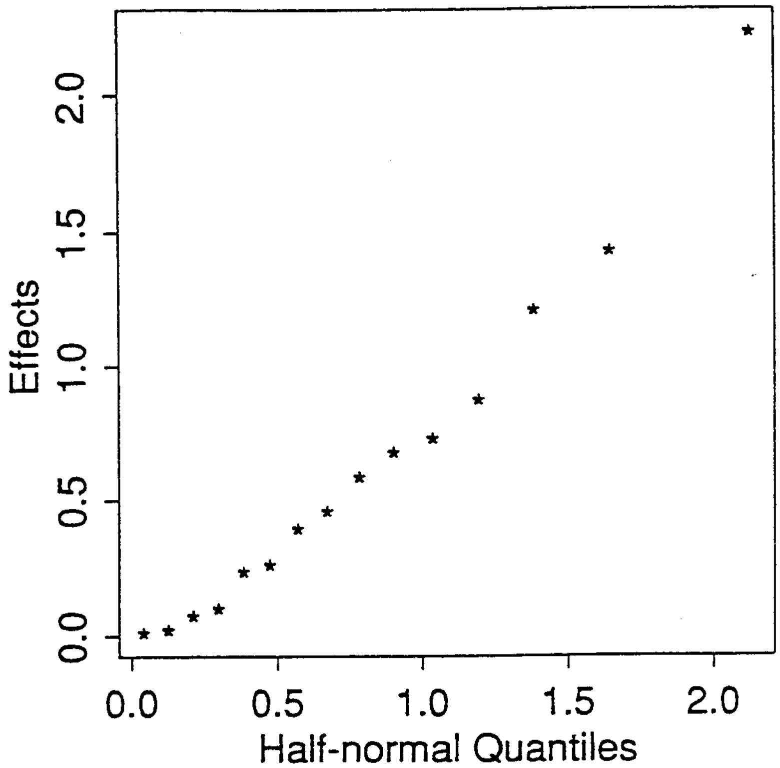

Figure 7. Half-normal plot of effects of log (variance) of studentized residuals. The 15 effects are the result of an analysis of variance fitting the model $\log(\hat{\sigma}_i^2) = A * B * C * D$. The $\hat{\sigma}_i^2$'s are computed using the residuals from fitting the model in (4.8).

Table 5. Effects of factors

	$\hat{\mu}_i$'s		$\hat{\beta}_i$'s		
	Cumulative power dissipation		Cumulative power dissipation	Number of fans	
Power dissipation	0	200		1	3
0	−0.11	9.91	0	0.84	0.96
40	16.09	26.11	200	3.23	1.13

We conclude that :

1. Shelf location has no effect on surface temperature.
2. Both the number of fans and the cumulative amount of power dissipated have dispersion effects, but there is interaction. When there is no power dissipated below the shelf of interest, the number of fans has no appreciable effect, and β_i is about one in both cases. When the cumulative amount of power dissipated is at its high level, increasing the number of fans from one to three reduces β_i from about three to nearly one. Since it is likely that there will be some power dissipated below the shelf, we should design the cabinet with three fans to reduce variability induced by potential changes in ambient temperature.

3. The amounts of power dissipated, both on the shelf of interest and cumulative power below the shelf, exhibit strong location effects. Increasing the amount of power dissipated on the shelf from 0 to 40 watts per circuit pack increases surface temperature by about 16 degrees. Similarly, increasing the cumulative power dissipated below the shelf from 0 to 200 watts per circuit pack increases surface temperature by about 10 degrees. These effects are approximately additive. These conclusions can now be combined with engineering knowledge to locate the high-power circuit packs and design the cabinet appropriately.

4.3. Discussion

- Often, engineering knowledge or other prior information will suggest the form of $g(\cdot)$. If not, and if k, the number of replications at each design setting, is not too small, we can plot the response against the noise variable at each setting and try to identify a suitable relationship. Transformations of one or both of the axes may be necessary to get a simple relationship. If we find a tentative functional form, we can use subsequent residual analyses to verify its validity.
- We also need to identify the link functions $h_0(\cdot)$ and $h_1(\cdot)$ in (3.2) and (3.3), and these are usually unknown in practice. If the data are extensive, it is possible to estimate the link function from the data and also test whether an assumed link function is adequate (see Pregibon (1980)). In industrial experiments with highly fractional designs, however, there is typically not enough information to estimate them from the data. Therefore, a priori knowledge or other considerations should be used, if at all possible, to select the appropriate link functions. For analyzing proportion data, for example, it is reasonable to assume a logit link function. In the absence of such knowledge, we recommend doing the analysis for several different reasonable link functions and choosing appropriate ones based on the criteria of parsimony and interpretability. The λ-plot discussed by Box (1988) will be a useful tool in this stage of the analysis. We note, however, that if the responses ($\hat{\beta}_i$'s and $\hat{\mu}_i$'s in this case) are all of the same order of magnitude, the conclusions from any link function will be qualitatively the same, so that the choice is not critical.
- As we discussed in the last section, a valid analysis of the $\hat{\beta}_i$'s should be based on weighted least-squares. If we are fitting a link function in (3.3) that is different from the identity function, we can compute the approximate variances of $h(\hat{\beta}_i)$ using the delta method and use the corresponding weights for doing weighted least squares. Similar comments apply to the analysis of $\hat{\mu}_i$'s.

- If the $\hat{\beta}_i$'s and $\hat{\mu}_i$'s in Step 2 are highly correlated, it may not be very meaningful to do model selection separately to identify the active location and dispersion effects. In this case, we recommend replacing Step 2 of the strategy by

Step 2'. Combine (3.1) with (3.2) and (3.3) to arrive at the global model (4.1), estimate the parameters $\boldsymbol{\alpha}$ and $\boldsymbol{\phi}$ directly and use a C_p-analysis based on overall prediction error to do model selection.

- The residuals are studentized in Step 4 to have (approximately) equal variance under the assumption that $\sigma_i^2 \equiv$ constant. When the σ_i^2's are unequal, these studentized residuals will not have equal variances. It is possible to develop a more refined analysis, based on an iterative procedure, that takes into account the unequal σ_i^2's. This is left for future work.
- Usually there is not enough information to identify the function $\nu(\cdot)$ in (3.8) empirically. In this case, we recommend making the usual assumption that this model is log-linear.
- In many situations, the conclusions from the statistical analysis will not be straightforward, and one has to resort to more formal means to determine the factor settings that achieve improved performance. This can be done by considering loss functions and distributions for the noise variables and determining the factor settings that minimize expected loss. See Welch et al. (1990) and Shoemaker et al. (1991) for more details associated with this formal approach.

5. Generalizations

5.1. Multiple noise variables

When there are several noise variables, it is natural to consider the following extension of (3.1). Suppose that, after a suitable transformation of the response, if necessary,

$$Y_{ij} = \mu_i + g_1(n_{1ij}; \beta_{1i}) + \cdots + g_r(n_{rij}; \beta_{ri}) + \varepsilon_{ij}, \tag{5.1}$$

where

$$h_0(\mu_i) = \boldsymbol{x}_i'\boldsymbol{\alpha}, \tag{5.2}$$

and

$$h_t(\beta_{ti}) = \boldsymbol{x}_i'\boldsymbol{\phi}_t,\ t = 1, \ldots, r. \tag{5.3}$$

We assume here that $k \geq r + 1$.

For $j \geq 1$, the ϕ_{tj}'s in (5.3) are measures of (generalized) interactions between the tth noise variable and design factors and indicate which design factors can be instrumental in reducing the effect of the tth noise variable on the response. If ϕ_{t0} is large and the ϕ_{tj}'s are all zero for $j \geq 1$ for some t, the effect of the tth noise variable cannot be reduced by robust parameter design and one may have to resort to other means such as tolerance design. Thus, model (5.1) allows one to study separately the effects of the individual noise variables and their relationship with the design factors. It is possible that some of the noise variables in (5.1) are controlled while others vary across the design settings. For example, in the thermal design experiment, measurements were made at three different locations – center, backplane, and faceplate. We can view "location" as a controlled noise variable and ambient temperature as uncontrolled, and analyze the surface temperatures at all three locations jointly using the model in (5.1).

One can extend the data-analysis strategy in Section 4 in an obvious way to handle multiple noise variables. In this case, however, there may be collinearity among the different (observed) noise variables, and it will be more meaningful to use Step 2′ and use an overall criterion for model selection. In the important special case where all the functions are linear, we can write the overall model as

$$Y_{ij} = \boldsymbol{x}_i'\boldsymbol{\alpha} + n_{1ij}\boldsymbol{x}_i'\boldsymbol{\phi}_1 + \cdots + n_{rij}\boldsymbol{x}_i'\boldsymbol{\phi}_r + \varepsilon_{ij}. \tag{5.4}$$

This can be expressed more compactly as

$$Y = \boldsymbol{Z}\boldsymbol{\theta} + \varepsilon \tag{5.5}$$

where $Y = [Y_{11}, \ldots, Y_{1k}, \ldots, Y_{n1}, \ldots, Y_{nk}]'$, ε is similarly defined, $\boldsymbol{\theta} = [\boldsymbol{\alpha}, \boldsymbol{\phi}_1, \ldots; \boldsymbol{\phi}_t]'$ and $\boldsymbol{Z}$ is the overall design matrix with elements specified by (5.4). This is a standard multiple regression problem with a structured design matrix. We can use standard techniques for diagnostics, model selection, and inference. One note of caution, however. The identified submodels are usually selected on predictive ability. If the overall regression matrix $\boldsymbol{Z}$ in (5.5) is highly collinear, there may be several different submodels that may perform equally well but that have different interpretations. So, one should exercise caution in interpreting the active location and dispersion effects in this situation.

5.2. Estimation with few or no replications

So far we have assumed that $k \geq r+1$, where k is the number of replications and r is the number of observed noise variables (or covariates). This may not always be the case in industrial experiments where it is important to keep the size of the experiment small. If there is effect sparsity, one may still be able to identify the active location and dispersion effects. We have to use different model selection techniques now instead of the C_p-plot based on all subset selection. There are potential identifiability problems here; and further work is needed to develop a careful iterative strategy, such as the one developed by Hamada and Wu (1991) for analyzing censored data. See also Box and Meyer (1986) for a strategy for estimating dispersion effects with no replication under the usual set-up.

5.3. Other situations

Sometimes the observed responses may be lifetimes which may be subject to censoring. It is possible to extend the techniques here to handle censored data. The analysis of censored data from highly fractional designs can be tricky, however (see Hamada and Wu (1991)).

Problems may also arise with the analysis of other nonstandard data such as count or proportion data. It is possible to handle these situations by applying transformations and analyzing the data using the strategy proposed here. An alternative that may sometimes be preferable is the use of extended quasi-likelihood models (see Nelder and Lee (1991)).

References

Box, G. E. P. (1988). Signal-to-noise ratios, performance criteria, and transformations (with discussion). *Technometrics* **30**, 1–40.

Box, G. E. P. and Meyer, R. D. (1986). Dispersion effects from fractional designs. *Technometrics* **28**, 19–27.

Hamada, M. and Wu, C. F. J. (1991). Analysis of censored data from highly fractionated experiments. *Technometrics* **33**, 25–38.

Léon, R. V., Shoemaker, A. C. and Kacker, R. N. (1987). Performance measures independent of adjustment: An explanation and extension of Taguchi's signal-to-noise ratios (with discussion). *Technometrics* **29**, 253–285.

Lucas, J. M. (1990). Achieving a robust process using response surface methodology. Technical Report, E. I. Du Pont de Nemours Company. Wilmington, DE.

Mallows, C. L. (1973). Some comments on C_p. *Technometrics* **15**, 661–675.

Nair, V. N. and Pregibon, D. (1986). A data-analysis strategy for quality engineering experiments. *AT&T Tech. J.* **65**, 73–84.

Nelder J. A. and Lee, Y. (1991). Generalized linear models for the analysis of Taguchi-type experiments. *Appl. Stochastic Models and Data Anal.* **7**, 107–120.

Phadke, M. S. (1989). *Quality Engineering Using Robust Design.* Prentice Hall.

Pregibon, D. (1980). Goodness of link tests for generalized linear models. *Appl. Statist.* **29**, 15–24.

Shoemaker, A. C., Tsui, K.-L. and Wu, C. F. J. (1991). Economical experimentation methods for robust design. *Technometrics* **33**, 415–427.

Taguchi, G. (1986). *Introduction to Quality Engineering.* Asian Productivity Organization, Tokyo.

Taguchi, G. (1987). *System of Experimental Design*, Vol. I and II. UNIPUB, New York.

Welch, W. J., Yu, T. K., Kang, S. M. and Sacks, J. (1990). Computer experiments for quality control by parameter design. *J. Quality Technology* **22**, 15–22.

AT&T Bell Laboratories, Murray Hill, NJ 07974–0636, U.S.A.

(Received February 1991; accepted February 1992)